QUANTIFIER ELIMINATION IN ORDERED ABELIAN GROUPS

RAF CLUCKERS AND IMMANUEL HALUPCZOK

ABSTRACT. We give a new proof of quantifier elimination in the theory of all ordered abelian groups in a suitable language. More precisely, this is only "quantifier elimination relative to ordered sets" in the following sense. Each definable set in the group is a union of a family of quantifier free definable sets, where the parameter of the family runs over a set definable (with quantifiers) in a sort which carries the structure of an ordered set with some additional unary predicates.

As a corollary, we find that all definable functions in ordered abelian groups are piecewise affine linear on finitely many definable pieces.

INTRODUCTION

Quantifier elimination is well known in some particular ordered abelian groups like $\mathbb{Q}$ and $\mathbb{Z}$. Somewhat less well known is that there also exists a quantifier elimination result for the theory of all ordered abelian groups. For sentences, this has already been proven by Gurevich [5] in 1964. Later, Gurevich and Schmitt enhanced this to treat arbitrary formulas ([3], [9]). The main goal of the present article is to introduce a new language L_{qe} with similar kind of quantifier elimination, which is more intuitive and hopefully more useful for applications.

As a corollary, we obtain that every definable function $f\colon G^n \to G$ in ordered abelian groups is piecewise linear, i.e., there exists a partition of G^n into finitely many definable sets such that the restriction of f to any of these sets is of the form $f(x_1,\ldots,x_n) = \frac{1}{s}(\sum_i r_i x_i + b)$ with $r_i, s \in \mathbb{Z}$ and $b \in G$. This result has been proven in the special case of groups of finite regular rank by Belegradek-Verbovskiy-Wagner [1] (using a version of quantifier elimination in this context from Weispfenning, [11]), but to our knowledge, it has yet not been written down in

full generality before. Our interest in this result came from valued fields. In the p-adics, definable maps can piecewise be approximated by fractional polynomials; see [2]. To get a similar result in valued fields with arbitrary value group, one necessary ingredient is piecewise linearity of definable maps in the value group.

Our quantifier elimination result could be deduced rather easily from the results of Gurevich and Schmitt. However, we discovered their results only after we had already written our own complete proof. We decided to include our proof in this article anyway to keep it self-contained and because both [3] and [9] are difficult to obtain. Moreover we are using a more modern formalism; in particular, we shall see that this kind of quantifier elimination naturally lives in many-sorted structures, which were not used by Gurevich and Schmitt.

From now on, we write "oag" for "ordered abelian group".

There is no really simple language in which oags have quantifier elimination; the main reason is that oags may have many convex definable subgroups, which come in several definable families. Parametrizing one such family with a suitable imaginary sort yields a uniform way to interpret an arbitrary ordered set in an appropriate oag. Since ordered sets have no good quantifier elimination language, the best one can hope for in oags is "quantifier elimination relative to ordered sets"; this is indeed we get.

Let us examine more closely what is needed in a quantifier elimination language. Recall that in the oag $\mathbb{Z}$, we have quantifier elimination in the Presburger language $L_{\mathrm{Pres}} := \{0, 1, +, <, \equiv_m\}$ (where $a \equiv_m b$ iff $a - b \in m\mathbb{Z}$). The same language also yields quantifier elimination in any fixed oag without (non-trivial) convex definable subgroup; in that case, 1 is defined to be the minimal positive element if this exists and $1 = 0$ otherwise. If G is a fixed group with finitely many convex definable subgroups H, then the quotients G/H are interpretable in G, and to get quantifier elimination, it is necessary (and sufficient) to have L_{Pres} not only on G, but also on all those quotients.

Now let us sketch the complete quantifier elimination language L_{qe}; it should allow for oags with infinite families of convex definable subgroups and moreover we want to work in the theory of all oags and not just in a fixed one. To treat infinite families of convex definable subgroups, we will add new sorts to L_{qe} (called "auxiliary sorts") with canonical parameters for some of them; let us write G_α for the group corresponding to the canonical parameter α. We will still need the Presburger language on all quotients G/G_α; roughly this will be formalized as follows: each quantifier free binary L_{Pres}-relation $x \diamond y$ (for $x, y \in G/G_\alpha$) becomes a ternary relation $\tilde{x} \diamond_\alpha \tilde{y}$ (for $\tilde{x}, \tilde{y} \in G$ and α in

an auxiliary sort) which holds iff the images of $\tilde{x}$ and $\tilde{y}$ in G/G_α satisfy $\diamond$. (For example, for each $m \in \mathbb{N}$, we have a relation $\tilde{x} \equiv_{m,\alpha} \tilde{y}$ which holds iff $\tilde{x} - \tilde{y} \in mG + G_\alpha$.)

Apart from that, three more things are needed in the language L_{qe}. On the auxiliary sorts, we have the order relation induced by inclusion of the corresponding subgroups and some unary predicates corresponding to certain properties of the groups G/G_α (which otherwise could not be expressed without quantifiers); moreover, we will need a variant of the congruence relation $\equiv_{m,\alpha}$ introduced above.

Our main result (Theorem 1.8) is that in L_{qe}, we have "quantifier elimination relative to the auxiliary sorts" in the following strong sense. Every definable subset in G is a union of a family of quantifier free definable sets, parametrized by an auxiliary set. This auxiliary set is defined by a formula which may use quantifiers, but it uses only the auxiliary part of L_{qe} (i.e., some ordered sets with unary predicates).

This kind of relative quantifier elimination might sound weak, but despite the fact that ordered sets have no good quantifier elimination, their model theory is well understood; see e.g. [8] or [7, ch. 12.f]. (This is also true for ordered sets with unary predicates, also called "colored chains".) Relative quantifier elimination allows to lift good model theoretic properties from ordered sets to oags; for example, Gurevich and Schmitt did this for NIP in [4]. Other results about oags may be deduced directly from relative quantifier elimination, without any knowledge of the auxiliary sorts at all; an example for this is our corollary about piecewise linearity of definable maps.

To prove relative quantifier elimination in L_{qe}, it is useful to simultaneously prove it in a second language L_{syn} which has certain good syntactic properties. These allow us to reduce relative quantifier elimination to eliminating a single existential quantifier of a formula which contains no other quantifiers, as one does it in usual quantifier elimination. This language L_{syn} is very close to the one used by Gurevich and Schmitt in their quantifier elimination results.

The article is organized as follows.

In Section 1, we present the main results: quantifier elimination in the languages L_{qe} and L_{syn} (Theorems 1.8 and 1.13) and piecewise linearity of definable functions (Corollary 1.10). We also state the general result on relative quantifier elimination in languages with good syntactic properties (Proposition 1.11). In this section, the languages are defined as quickly as possible, postponing explanations to the next section. At the end of the section, we explain the relation between L_{syn} and the language used by Schmitt.

In Section 2, we prove some first basic properties of the languages, which also yields some motivation. Then we show how to translate between L_{syn} and L_{qe}, allowing us to switch freely between those languages while doing quantifier elimination.

Section 3 contains the main proofs. First, we prove Proposition 1.11. Then we do the actual elimination of one existential quantifier; this is done in the language L_{qe}. The whole proof is constructive, so it can be turned into an algorithm for quantifier elimination.

Section 4 contains some examples illustrating the language L_{qe}; in particular, they show how arbitrary ordered sets can be interpreted in oags.

0.1. **Acknowledgment.** We are very grateful to Françoise Delon for a lot of interesting discussions and for several valuable concrete suggestions concerning this article.

1. The results

1.1. **Generalities and basic notation.** We use the convention that $0 \notin \mathbb{N}$, and we write $\mathbb{N}_0$ for $\mathbb{N} \cup \{0\}$ and $\mathbb{P}$ for the set of primes.

In the whole article, $(G, +, <)$ will denote an ordered abelian group ("oag"), that is, a group with a total order which is compatible with the group operation: $a < b$ iff $a + c < b + c$ for all $a, b, c \in G$. It is easy to see that such a group is always torsion free. Such groups appear naturally, for example, as valuation group of (Krull) valued fields. An oag is called *discrete*, if it has a minimal positive element and *dense* otherwise.

We write $L_{\mathrm{oag}} = \{0, +, <\}$ for the language of oags and unless stated otherwise, we always work in the theory of all oags.

For $a \in G$, we write $\langle a \rangle^{\mathrm{conv}}$ for the smallest convex subgroup of G containing a; for $a, b \in G$ and $m \in \mathbb{N}$, $a \equiv_m b$ means that a and b are congruent modulo m in the sense that $a - b \in mG$.

1.2. **A language for quantifier elimination.** We now give a precise definition of the quantifier elimination language L_{qe}; motivation and and additional explanations will be given in Section 2. An introduction to L_{qe} with much more motivation and examples can be found in [6]. Note that all of L_{qe} will be L_{oag}-definable (where new sorts in L_{qe} are considered as imaginary sorts of L_{oag}).

We start by introducing the new sorts of L_{qe}: sorts with canonical parameters for some definable families of convex subgroups. These new sorts will be called *auxiliary sorts*; in contrast, the sort of the ordered abelian group itself will be called the *main sort*.

For each positive integer n, we consider three families of convex definable subgroups, parametrized by sorts which we denote by $\mathcal{S}_n$, $\mathcal{T}_n$, and $\mathcal{T}_n^+$. Although in L_{qe} we will have these sorts only for n prime, it is useful to define them for all n. Examples illustrating the following definition are given in Section 4.

Definition 1.1.   (1) For $n \in \mathbb{N}$ and $a \in G \setminus nG$, let H_a be the largest convex subgroup of G such that $a \notin H_a + nG$; set $H_a = \{0\}$ if $a \in nG$. Define $\mathcal{S}_n := G/\sim$, with $a \sim a'$ iff $H_a = H_{a'}$, and let $\mathfrak{s}_n \colon G \twoheadrightarrow \mathcal{S}_n$ be the canonical map. For $\alpha = \mathfrak{s}_n(a) \in \mathcal{S}_n$, define $G_\alpha := H_a$.
(2) For $n \in \mathbb{N}$ and $b \in G$, set $H_b' := \bigcup_{\alpha \in \mathcal{S}_n, b \notin G_\alpha} G_\alpha$, where the union over the empty set is $\{0\}$. Define $\mathcal{T}_n := G/\sim$, with $b \sim b'$ iff $H_b' = H_{b'}'$, and let $\mathfrak{t}_n \colon G \twoheadrightarrow \mathcal{T}_n$ be the canonical map. For $\alpha = \mathfrak{t}_n(b) \in \mathcal{T}_n$, define $G_\alpha := H_b'$.
(3) For $n \in \mathbb{N}$ and $\beta \in \mathcal{T}_n$, define $G_{\beta+} := \bigcap_{\alpha \in \mathcal{S}_n, G_\alpha \supsetneq G_\beta} G_\alpha$, where the intersection over the empty set is G. Here, we view the index $\beta+$ as being an element of a copy of $\mathcal{T}_n$ which we denote by $\mathcal{T}_n^+$.
(4) Define a total preorder on $\dot\bigcup_{n \in \mathbb{N}} (\mathcal{S}_n \mathbin{\dot\cup} \mathcal{T}_n \mathbin{\dot\cup} \mathcal{T}_n^+)$ by $\alpha \le \alpha'$ iff $G_\alpha \subseteq G_{\alpha'}$. Write $\alpha \asymp \alpha'$ if $G_\alpha = G_{\alpha'}$.

Definability (in L_{oag}) of the groups G_α, $\alpha \in \mathcal{S}_n$ is proven in Lemma 2.1; once this is done, it is clear that the new sorts are imaginary sorts of L_{oag} and that all of the above is definable.

Remark 1.2. If $b \ne 0$, then we have $G_{\mathfrak{t}_n(b)+} = \bigcap_{\alpha \in \mathcal{S}_n, b \in G_\alpha} G_\alpha$; in particular, $G_{\mathfrak{t}_n(b)+}$ is strictly bigger than $G_{\mathfrak{t}_n(b)}$, since $b \notin G_{\mathfrak{t}_n(b)}$. (However, we might have $G_{\mathfrak{t}_n(0)+} = \{0\}$.)

Fix α in any of the auxiliary sorts. Recall that for each quantifier free L_{Pres}-definable relation on G/G_α, we want the corresponding relation on G to be quantifier free definable in L_{qe}. If G/G_α is dense, then it suffices to put preimages of the relations $=, <, \equiv_m$ into L_{qe} (interpreted as ternary relations, where α is the third operand). However, if G/G_α has a minimal positive element, then we need L_{qe}-predicates for preimages of L_{Pres}-relations defined using this element. We introduce the following notation for these predicates.

Definition 1.3. Suppose that $\alpha \in \mathcal{S}_n \cup \mathcal{T}_n \cup \mathcal{T}_n^+$ for some $n \in \mathbb{N}$ and that $\pi \colon G \twoheadrightarrow G/G_\alpha$ is the canonical projection. For $\diamond \in \{=, <, >, \le, \ge, \equiv_m\}$, write $x \diamond_\alpha y$ if $\pi(x) \diamond \pi(y)$ holds in G/G_α.

For $k \in \mathbb{Z}$, write k_α for k times the minimal positive element of G/G_α if G/G_α is discrete and set $k_\alpha := 0 \in G/G_\alpha$ otherwise. Write $x \diamond_\alpha y + k_\alpha$ for $\pi(x) \diamond \pi(y) + k_\alpha$.

Note that $x \equiv_{m,\alpha} y$ holds iff $x - y \in G_\alpha + mG$. We will need one additional kind of predicates which is similar, but where G_α is replaced by a group which looks rather technical. For definability of that group and for more explanations, see Section 2.2.

Definition 1.4. For $n, m, m' \in \mathbb{N}$ and $\alpha \in \mathcal{S}_n \cup \mathcal{T}_n \cup \mathcal{T}_n^+$, set

$$G_\alpha^{[m']} := \bigcap_{\substack{H \subseteq G \text{ convex subgroup} \\ H \supsetneq G_\alpha}} (H + m'G);$$

write $x \equiv_{m,\alpha}^{[m']} y$ iff $x - y \in G_\alpha^{[m']} + mG$.

A separate notation for $x - y \in G_\alpha^{[m']}$ is not needed, since $G_\alpha^{[m']} = G_\alpha^{[m']} + m'G$.

Finally, in L_{qe} we will need a few unary predicates on the auxiliary sorts: one saying whether the group G/G_α is discrete, and some predicates specifying the cardinalities of certain quotients of two groups of the form $G_\alpha + pG$ or $G_\alpha^{[p^s]} + pG$. Since pG is contained in the denominator of those quotients, they are $\mathbb{F}_p$ vector spaces, and specifying the cardinality is equivalent to specifying the dimension over $\mathbb{F}_p$.

Here is the complete definition of L_{qe}:

Definition 1.5. The language L_{qe} consists of the following:
- The main sort G with the constant 0, the binary function $+$, and the unary function $-$.
- For each $p \in \mathbb{P}$, the auxiliary sorts $\mathcal{S}_p$, $\mathcal{T}_p$ and $\mathcal{T}_p^+$ from Definition 1.1.
- For each $p, p' \in \mathbb{P}$: binary relations "$\alpha \le \alpha'$" on $(\mathcal{S}_p \mathbin{\dot{\cup}} \mathcal{T}_p \mathbin{\dot{\cup}} \mathcal{T}_p^+) \times (\mathcal{S}_{p'} \mathbin{\dot{\cup}} \mathcal{T}_{p'} \mathbin{\dot{\cup}} \mathcal{T}_{p'}^+)$, defined by $G_\alpha \subseteq G_{\alpha'}$. (For each p, p', these are nine relations.)
- Predicates for the relations $x_1 \diamond_\alpha x_2 + k_\alpha$ from Definition 1.3, where $\diamond \in \{=, <, \equiv_m\}$, $k \in \mathbb{Z}$, $m \in \mathbb{N}$, and where α may be from any of the sorts $\mathcal{S}_p$, $\mathcal{T}_p$ and $\mathcal{T}_p^+$. (These are ternary relations on $G \times G \times \mathcal{S}_p$, $G \times G \times \mathcal{T}_p$, and $G \times G \times \mathcal{T}_p^+$.)
- For each $p \in \mathbb{P}$ and each $m, m' \in \mathbb{N}$, the ternary relation $x \equiv_{m,\alpha}^{[m']} y$ on $G \times G \times \mathcal{S}_p$.
- For each $p \in \mathbb{P}$, a predicate $\mathrm{discr}(\alpha)$ on $\mathcal{S}_p$ which holds iff G/G_α is discrete.
- For each $p \in \mathbb{P}$, each $s \in \mathbb{N}$, and each $\ell \in \mathbb{N}_0$, two predicates on $\mathcal{S}_p$ defining the sets

$$\{\alpha \in \mathcal{S}_p \mid \dim_{\mathbb{F}_p}(G_\alpha^{[p^s]} + pG)/(G_\alpha^{[p^{s+1}]} + pG) = \ell\} \quad \text{and}$$

$$\{\alpha \in \mathcal{S}_p \mid \dim_{\mathbb{F}_p}(G_\alpha^{[p^s]} + pG)/(G_\alpha + pG) = \ell\}.$$

Notation 1.6. We write $\mathcal{M} := \{G\}$ for the main sort and $\mathcal{A} := \{\mathcal{S}_p, \mathcal{T}_p, \mathcal{T}_p^+ \mid p \in \mathbb{P}\}$ for the collection of auxiliary sorts. By abuse of notation, we will also write $\mathcal{A}$ for the union of the auxiliary sorts. We will write that a formula is "$\mathcal{M}$-qf" if it does not contain any quantifier running over a main sort variable.

The usual predicates $<$ and $\equiv_m$ on G are $\mathcal{M}$-qf L_{qe}-definable: they are equivalent to $<_{\alpha_0}$ and $\equiv_{m,\alpha_0}$, where α_0 is the minimal element of, say, $\mathcal{S}_2$. The canonical map $\mathcal{T}_p \to \mathcal{T}_p^+, \alpha \mapsto \alpha+$ is easily $\mathcal{M}$-qf definable from the preorder on $\mathcal{T}_p \mathbin{\dot{\cup}} \mathcal{T}_p^+$ using Remark 1.2. We will later see $\mathcal{M}$-qf definability of the canonical maps $\mathfrak{s}_p$, $\mathfrak{t}_p$ (Lemma 2.8) and of the analogues on $\mathcal{T}_p$ and $\mathcal{T}_p^+$ of the discreteness and dimension predicates (Lemmas 2.11 and 2.10). Moreover, Lemmas 2.2 and 2.3 will show how to get along without having $\mathcal{S}_n$, $\mathcal{T}_n$, $\mathcal{T}_n^+$, $\mathfrak{s}_n$, and $\mathfrak{t}_n$ for arbitrary n.

Note that although $\mathcal{T}_p$ and $\mathcal{T}_p^+$ are in definable bijection, identifying them would make the language pretty messy, in particular because the preorder on $\dot{\bigcup}_p (\mathcal{S}_p \mathbin{\dot{\cup}} \mathcal{T}_p)$ is *not* enough to define the preorder on the whole of $\mathcal{A}$ in an $\mathcal{M}$-qf way.

As announced, our main result is "quantifier elimination relative to the auxiliary sorts", which is more than just elimination of main sort quantifiers. Now let us make this precise; we first need a definition.

Definition 1.7. Suppose that L is any language, T is an L-theory, $\mathcal{M} \mathbin{\dot{\cup}} \mathcal{A}$ is a partition of the sorts of L, and $\phi(\bar{x}, \bar{\eta})$ is an L-formula, where $\bar{x}$ are $\mathcal{M}$-variables and $\bar{\eta}$ are $\mathcal{A}$-variables. We say that $\phi(\bar{x}, \bar{\eta})$ is in *family union form* if it is of the form

$$\phi(\bar{x}, \bar{\eta}) = \bigvee_{i=1}^{k} \exists \bar{\theta} \left(\xi_i(\bar{\eta}, \bar{\theta}) \wedge \psi_i(\bar{x}, \bar{\theta}) \right)$$

where $\bar{\theta}$ are $\mathcal{A}$-variables, the formulas $\xi_i(\bar{\eta}, \bar{\theta})$ live purely in the sorts $\mathcal{A}$, each $\psi_i(\bar{x}, \bar{\theta})$ is a conjunction of literals (i.e., of atoms and negated atoms), and T implies that the $L(\mathcal{A})$-formulas $\{\xi_i(\bar{\eta}, \bar{\alpha}) \wedge \psi_i(\bar{x}, \bar{\alpha}) \mid 1 \leq i \leq k, \bar{\alpha} \in \mathcal{A}\}$ are pairwise inconsistent.

Theorem 1.8. *In the theory of ordered abelian groups, each L_{qe}-formula is equivalent to an L_{qe}-formula in family union form.*

Remark 1.9. In L_{qe}, the formulas $\psi_i(\bar{x}, \bar{\theta})$ appearing in the family union form are very simple. Without loss, each atom involves the main sort, i.e., it is of the form $t(\bar{x}) \diamond_{\theta_\nu} t'(\bar{x}) + k_{\theta_\nu}$ where $t(\bar{x}), t'(\bar{x})$ are $\mathbb{Z}$-linear combinations, $\diamond \in \{=, <, \equiv_m, \equiv_m^{[m']}\}$, θ_ν is one of the entries of $\bar{\theta}$, $k \in \mathbb{Z}$, and $m, m' \in \mathbb{N}$ (where $k = 0$ if $\diamond$ is $\equiv_m^{[m']}$). Moreover, "$=$"-literals can be expressed using "$<$" and "$>$" instead. Now the

inequalities of ψ_i define a convex polyhedron, and the remaining literals $(\equiv_m, \not\equiv_m, \equiv_m^{[m']}, \not\equiv_m^{[m']})$ are "congruence conditions" in the sense that each of them defines a set which consists of entire cosets of mG (possibly for several different $m \in \mathbb{N}$). From this point of view, such sets are very similar to sets definable in L_{Pres} by a conjunction of literals (which are also intersections of polyhedra with congruence conditions).

1.3. **Definable functions are piecewise linear.** Using the above quantifier elimination theorem, it is easy to prove that definable functions from G^n to G are piecewise linear. More precisely:

Corollary 1.10. *For any function $f \colon G^n \to G$ which is L_{oag}-definable with parameters from a set B, there exists a partition of G^n into finitely many B-definable sets such that on each such set A, f is linear: there exist $r_1, \ldots r_n, s \in \mathbb{Z}$ with $s \neq 0$ and $b \in \mathrm{dcl}(B)$ such that for any $\bar{a} \in A$, we have $f(a_1, \ldots, a_n) = \frac{1}{s}(\sum_i r_i a_i + b)$.*

Let us prove this right away, since it illustrates nicely how Theorem 1.8 can be applied.

Proof. Let $\phi(\bar{x}, y)$ be an $L_{\text{qe}}(B)$-formula in family union form defining the graph of f, let $\bar{a} \in G^n$ be a tuple, set $c := f(\bar{a})$, and consider $\phi(\bar{a}, y) \in L_{\text{qe}}(B \cup \bar{a})$, which defines the single element set $\{c\}$. (We do not write the parameters from B explicitly.) Using a case distinction, we may suppose that the family union form of $\phi(\bar{a}, y)$ consists of a single family:

$$\phi(\bar{a}, y) = \exists \bar{\theta} \left(\xi(\bar{\theta}) \wedge \psi(\bar{a}, y, \bar{\theta}) \right).$$

Let $\bar{\beta}$ be the (unique) tuple of $\mathcal{A}$ such that $G \models \psi(\bar{a}, c, \bar{\beta})$.

As in Remark 1.9, we may assume that $\psi(\bar{a}, y, \bar{\beta})$ uses no "=". Moreover, we may choose an $m_0 \in \mathbb{N}$ such that all congruence conditions of $\psi(\bar{a}, y, \bar{\beta})$ together define a union of cosets of $m_0 G$.

Using further case distinctions (which are definable in $\bar{a}$), we can assume: all literals of $\psi(\bar{a}, y, \bar{\beta})$ involve y and among these literals, there is at most one lower and one upper bound on y.

There has to be a lower bound; otherwise, for $d \in G$ with $d > 0$, the element $c - m_0 d$ would also satisfy $\psi(\bar{a}, y, \bar{\beta})$. We may suppose that the lower bound is of the form $ry \vartriangleright_\alpha t(\bar{a}) + k_\alpha$, where $\vartriangleright \in \{>, \geq\}$, $\alpha \in \mathcal{A}$, and where t is a main sort term, i.e., a $\mathbb{Z}$-linear combination of entries of $\bar{a}$ plus an element of $\mathrm{dcl}(B)$. If $G_\alpha \supsetneq \{0\}$, then again $c - m_0 d$ satisfies $\psi(\bar{a}, y, \bar{\beta})$ if we take $d \in G_\alpha, d > 0$; hence $G_\alpha = \{0\}$. In particular, k_α can be seen as an element of G (and not just as a notation). From this point of view, we have $k_\alpha \in \mathrm{dcl}(\emptyset)$, so without loss, the lower bound is of the form $ry \vartriangleright_\alpha t(\bar{a})$.

Since c is unique satisfying $\psi(\bar{a}, y, \bar{\beta})$, it must be the minimal element satisfying $ry \rhd_\alpha t(\bar{a})$ and the congruence conditions. Such a minimum can only exist if $G_\alpha = \{0\}$. If G is dense, then $m_0 G$ is dense in G, so a minimum has to be equal to the lower bound; thus ψ is equivalent to $ry = t(\bar{a})$ and we are done. If G is discrete, then we do a case distinction on the difference $d := rc - t(\bar{a})$. This difference can be at most $rm_0 + 1$ (otherwise $c - m_0$ would also satisfy $\psi(\bar{a}, y, \bar{\beta})$), so there are only finitely many cases. Fixing d is a definable condition on $\bar{a}$ and for fixed d, ψ is equivalent to $rc = t(\bar{a}) + d$, which again is linear. $\qquad\square$

1.4. **A language with good syntactic properties.** For usual quantifier elimination, it suffices to prove that the quantifier of $\exists x \, \psi(x)$ can be eliminated when $\psi(x)$ is quantifier free. This does not work for relative quantifier elimination: neither if we only try to get rid of $\mathcal{M}$-quantifiers (then ψ can contain $\mathcal{A}$-quantifiers, which can make it pretty complicated), nor if we want to get a formula in family union form (in that case, the main difficulty turns out to be that it is not clear whether formulas in family union form are closed under negation). The following general result allows us to do such a reasoning anyway under some syntactic assumptions about the language.

Proposition 1.11. *Let L be a language and let $\mathcal{M} \mathbin{\dot{\cup}} \mathcal{A}$ be a partition of the sorts of L. Suppose that the only symbols in L connecting $\mathcal{M}$ and $\mathcal{A}$ are functions from (products of) $\mathcal{M}$-sorts to $\mathcal{A}$-sorts. Let T be an L-theory.*

Consider a formula of the form $\exists x \, \psi(x, \bar{y}, \bar{\eta})$ where $x, \bar{y}$ are $\mathcal{M}$-variables, $\bar{\eta}$ are $\mathcal{A}$-variables and ψ is quantifier free. Suppose that modulo T, any such formula is equivalent to a formula without $\mathcal{M}$-quantifiers.

Then modulo T, any L-formula is equivalent to an L-formula in family union form.

Note that the proposition does not require us to bring $\exists x \, \psi(x, \bar{y}, \bar{\eta})$ into family union form; no $\mathcal{M}$-quantifiers is enough.

To be able to apply this result to ordered abelian groups, we introduce a second language L_{syn} which has the required property: all L_{qe}-predicates connecting $\mathcal{M}$ and $\mathcal{A}$ will be replaced by some predicates on $\mathcal{M}$ and some functions from $\mathcal{M}$ to $\mathcal{A}$. Let me start by explaining the idea of how this can be done; a complete proof that L_{syn} is as strong as L_{qe} will be given in Section 2.5.

The L_{qe}-predicates we have to get rid of are $x_1 \diamond_\eta x_2 + k_\eta$ for the various $\diamond$. First consider $x_1 =_\eta x_2$. Since for fixed x_1 and x_2, $x_1 =_\eta x_2$ holds if and only if η is bigger than a certain bound depending only on

$x_1 - x_2$, we can replace the predicate $x_1 =_\eta x_2$ by the function from G to $\mathcal{A}$ which returns this bound. In the case $\eta \in \mathcal{S}_p$, we already defined exactly this function: it is the canonical map $\mathfrak{t}_p \colon G \twoheadrightarrow \mathcal{T}_p$; for $\eta \in \mathcal{T}_p \mathbin{\dot{\cup}} \mathcal{T}_p^+$, one verifies that $\mathfrak{t}_p$ still works.

A similar idea allows to express the predicates $x_1 \equiv_{p^r,\eta} x_2$ using the canonical maps $\mathfrak{s}_{p^r}$ (for $p \in \mathbb{P}$ and $r \in \mathbb{N}$). In principle, these maps go to $\mathcal{S}_{p^r}$ which are not sorts of L_{syn}, but we will see in Lemma 2.2 that $\mathcal{S}_{p^r}$ and $\mathcal{S}_p$ can be identified.

What is missing now is a way to deal with the predicates $x_1 \diamond_\alpha x_2 + k_\alpha$ when $k \neq 0$ (for $\diamond \in \{=, \equiv_m\}$) and with $x \equiv^{[m']}_{m,\alpha} y$. (The inequalities $<_\alpha$ are no problem.) These predicates will essentially be replaced by their union over all α. We will see in Section 2.4 how the L_{qe}-predicates can be reconstructed from this.

Here is the complete definition of the language L_{syn}:

Definition 1.12. The language L_{syn} consists of the following:

- The main sort G with $0, +, -, <$, and $\equiv_m$ (for $m \in \mathbb{N}$).
- As in L_{qe}, the auxiliary sorts $\mathcal{S}_p$, $\mathcal{T}_p$ and $\mathcal{T}_p^+$ with the binary relations "$\alpha \leq \alpha'$" on $(\mathcal{S}_p \mathbin{\dot{\cup}} \mathcal{T}_p \mathbin{\dot{\cup}} \mathcal{T}_p^+) \times (\mathcal{S}_{p'} \mathbin{\dot{\cup}} \mathcal{T}_{p'} \mathbin{\dot{\cup}} \mathcal{T}_{p'}^+)$, and on $\mathcal{S}_p$ the unary predicates $\mathrm{discr}(\alpha)$, $\dim_{\mathbb{F}_p}(G_\alpha^{[p^s]} + pG)/(G_\alpha^{[p^{s+1}]} + pG) = \ell$, and $\dim_{\mathbb{F}_p}(G_\alpha^{[p^s]} + pG)/(G_\alpha + pG) = \ell$.
- For each $p \in \mathbb{P}$ (and each $r \in \mathbb{N}$): the canonical maps $\mathfrak{s}_{p^r} \colon G \twoheadrightarrow \mathcal{S}_p$ and $\mathfrak{t}_p \colon G \twoheadrightarrow \mathcal{T}_p$ from Definition 1.1, where $\mathcal{S}_{p^r}$ is identified with $\mathcal{S}_p$ using Lemma 2.2.
- For each $k \in \mathbb{Z} \setminus \{0\}$: a unary predicate "$x =_\bullet k_\bullet$" on G defined by: there exists a convex subgroup $H \subseteq G$ such that G/H is discrete and the image of x in G/H is k times the smallest positive element of G/H; see Section 2.4 for details, in particular for definability.
- For each $m \in \mathbb{N}$ and each $k \in \{1, \dots, m-1\}$: a unary predicate "$x \equiv_{m,\bullet} k_\bullet$" on G defined by: there exists a convex subgroup $H \subseteq G$ such that G/H is discrete and the image of x in G/H is congruent modulo m to k times the minimal positive element of G/H; again see Section 2.4 for details.
- For each $p \in \mathbb{P}$ and each $r, s \in \mathbb{N}$ with $s \geq r$: a unary predicate $D_{p^r}^{[p^s]}(x)$ on G for: there exists an $\alpha \in \mathcal{S}_p$ such that x lies in $G_\alpha^{[p^s]} + p^r G$, but not in $G_\alpha + p^r G$.

In this language, relative quantifier elimination will simply be the conclusion of Proposition 1.11:

Theorem 1.13. *In the theory of ordered abelian groups, each L_{syn}-formula is equivalent to an L_{syn}-formula in family union form.*

We will deduce Theorem 1.8 from this one by translating the $\mathcal{M}$-qf L_{syn}-formula back into L_{qe}. This will be done at the end of Section 2.5.

1.5. **Comparison to Gurevich and Schmitt.** Theorem 1.13 is very similar to the quantifier elimination results of Gurevich and Schmitt; here we give a little translation table between our language L_{syn} and the one used in Schmitt's habilitation thesis [9]. The quantifier elimination result of [9] (Lemma 4.3, Theorem 4.4) is also described in the introduction of [10] (Theorem 1.7).

Schmitt does not distinguish between the sorts $\mathcal{S}_n$, $\mathcal{T}_n$, and $\mathcal{T}_n^+$; instead, for each $n \in \mathbb{N}$ he works with a single sort $\mathrm{Sp}_n(G) := (\mathcal{S}_n \,\dot{\cup}\, \mathcal{T}_n \,\dot{\cup}\, \mathcal{T}_n^+)/\asymp$ (the "n-spine of G"), with predicates for $\mathcal{S}_n$ and $\mathcal{T}_n$. (More precisely, Schmitt does not really use a multi-sorted structure, but this is what his formulation boils down to.)

When eliminating the $\mathcal{M}$-quantifiers of a given formula ϕ, instead of using several sorts $\mathrm{Sp}_p(G)$ for primes p, he uses only one single sort $\mathrm{Sp}_n(G)$ for $n \in \mathbb{N}$.

Instead of our dimension predicates, Schmitt has predicates for the *Szmielew-invariants* of $G_{p^r,\alpha}^{[p^r]}/G_{p^r,\alpha}$ (see Definition on page 5 of [9]). At first sight, it seems that the number of Szmielew-invariants is bigger than the number dimensions for which we introduced predicates (for each α, the set of Szmielew-invariants is parametrized by two natural numbers, whereas we consider only two families of dimensions parametrized by a single natural number), but a little computation shows that many of the Szmielew-invariants are always equal (and equal to our dimensions).

Finally, on the main sort, Schmitt has slightly different predicates than our $x =_\bullet k_\bullet$, $x \equiv_{m,\bullet} k$ and $D_{p^r}^{[p^s]}$.

2. Details of the languages

2.1. **The families of convex definable subgroups G_α.** In Definition 1.1, we introduced the families of convex groups G_α, but we still had to verify that they are definable in the case $\alpha \in \mathcal{S}_n$.

Lemma 2.1. *Fix $n \in \mathbb{N}$. For $a \in G$, the group $G_{\mathfrak{s}_n(a)}$ is definable uniformly in a.*

Proof. We may suppose $a \notin nG$. In that case, $G_{\mathfrak{s}_n(a)}$ consists of those elements $b \in G$ such that $a \notin \langle b \rangle^{\mathrm{conv}} + nG$. The group $\langle b \rangle^{\mathrm{conv}}$ is not definable in general, but we have $\langle b \rangle^{\mathrm{conv}} + nG = [0, n|b|] + nG$, which is definable; here, $|b|$ denotes the absolute value of b. $\qquad\square$

We defined the sorts $\mathcal{S}_n$, $\mathcal{T}_n$ and $\mathcal{T}_n^+$ for arbitrary n, but in our languages, we only have them for n prime. The following two lemmas will allow us to reduce any usage of these sorts to the prime cases. In particular, we show that $\mathcal{S}_{p^r}$ can be identified with $\mathcal{S}_p$, as required in the definition of L_{syn}.

We use the notation "$p^r \mid\mid n$" from number theory which means that p is a prime divisor of n and that p^r is the maximal power of p dividing n.

Lemma 2.2. *Let $n \in \mathbb{N}$.*

(1) *We have the following equality of sets of convex subgroups of G:*

$$\{G_\alpha \mid \alpha \in \mathcal{S}_n\} = \bigcup_{p \in \mathbb{P}, p \mid n} \{G_\alpha \mid \alpha \in \mathcal{S}_p\}.$$

In particular, there is a (unique, definable) bijection $\mathcal{S}_{p^r} \to \mathcal{S}_p$ which is compatible with $\alpha \mapsto G_\alpha$.

(2) *For any $a \in G$, we have*

$$G_{\mathfrak{s}_n(a)} = \bigcup_{p^r \mid\mid n} G_{\mathfrak{s}_{p^r}(a)}.$$

In particular, $\mathfrak{s}_n(a) \asymp \max_{p^r \mid\mid n} \mathfrak{s}_{p^r}(a)$.

Proof. We start with (1) "$\supseteq$"; more precisely, for $m \mid n$, we prove $\{G_\alpha \mid \alpha \in \mathcal{S}_n\} \supseteq \{G_\alpha \mid \alpha \in \mathcal{S}_m\}$. Consider $G_\alpha \neq \{0\}$ in the right hand set and choose $a \in G \backslash mG$ with $\alpha = \mathfrak{s}_m(a)$. Recall that G_α is the largest convex subgroup of G with $a \notin G_\alpha + mG$. For any convex subgroup $H \in G$, we have $a \in H + mG$ if and only if $a' := \frac{n}{m}a \in H + nG$; hence $G_\alpha = G_{\mathfrak{s}_n(a')}$.

Next, we prove (2). The inclusion "$\supseteq$" is clear. For "$\subseteq$", we may suppose that $a \in G \setminus nG$. By the Chinese remainder theorem, we have $G_{\mathfrak{s}_n(a)} + nG = \bigcap_{p^r \mid\mid n}(G_{\mathfrak{s}_n(a)} + p^r G)$, so $a \notin G_{\mathfrak{s}_n(a)} + nG$ implies $a \notin G_{\mathfrak{s}_n(a)} + p^r G$ for some $p \mid n$. This in turn implies $G_{\mathfrak{s}_n(a)} \subseteq G_{\mathfrak{s}_{p^r}(a)}$.

Finally, we prove (1) "$\subseteq$". By (2), we have $\{G_\alpha \mid \alpha \in \mathcal{S}_n\} \subseteq \{G_\alpha \mid \alpha \in \bigcup_{p^r \mid\mid n} \mathcal{S}_{p^r}\}$, so it suffices to do the case where $n = p^r$. Suppose that $\alpha = \mathfrak{s}_{p^r}(a)$ for some $a \in G \setminus p^r G$ and consider the group G_α from the left hand set of (1). Let $s \in \mathbb{N}$ be maximal with $a \in G_\alpha + p^s G$; by assumption $s < r$. Write $a = b + p^s a'$ for $b \in G_\alpha$ and $a' \in G$. Then $a' \notin G_\alpha + pG$, since otherwise $b + p^s a' \in G_\alpha + p^{s+1}G$. On the other hand for any convex subgroup H strictly larger than G_α, we have $b + p^s a' = a \in H + p^r G \subseteq H + p^{s+1}G$, so $p^s a' \in H + p^{s+1}G$, so $a' \in H + pG$. Hence $G_{\mathfrak{s}_p(a')} = G_\alpha$. $\qquad\square$

Lemma 2.3. *For any $n \in \mathbb{N}$ and any $a \in G$, we have*

$$G_{\mathfrak{t}_n(a)} = \bigcup_{p \in \mathbb{P}, p \mid n} G_{\mathfrak{t}_p(a)} \quad and \quad G_{\mathfrak{t}_n(a)+} = \bigcap_{p \in \mathbb{P}, p \mid n} G_{\mathfrak{t}_p(a)+}.$$

In particular, $\mathfrak{t}_n(a) \asymp \max_{p \in \mathbb{P}, p \mid n} \mathfrak{t}_p(a)$ and $\mathfrak{t}_n(a)+ \asymp \min_{p \in \mathbb{P}, p \mid n}(\mathfrak{t}_p(a)+)$.

Proof. This follows directly from Lemma 2.2 (1), where for $G_{\mathfrak{t}_n(a)+}$ we use Remark 1.2. $\qquad\square$

2.2. **Congruence conditions and expressing $\mathfrak{s}_n$ and $\mathfrak{t}_n$ in L_{qe}.** In Definition 1.4, we introduced the group $G_\alpha^{[n]} = \bigcap_{H \supsetneq G_\alpha}(H + nG)$ for $n \in \mathbb{N}$ and $\alpha \in \mathcal{A}$. The point is that $G_\alpha^{[n]}$ might be strictly bigger than $(\bigcap_{H \supsetneq G_\alpha} H) + nG$, and in general, it is not of the form $H_0 + nG$ for any convex subgroup $H_0 \subseteq G$ (see example in Section 4.3). We will need these groups to express the L_{syn}-function $\mathfrak{s}_n$ in L_{qe} without $\mathcal{M}$ quantifiers; this will be done at the end of this section.

The following lemma gives an equivalent definition of $G_\alpha^{[n]}$ (using not all convex subgroups of G) which in particular shows that it is definable.

Lemma 2.4. *Let $n \in \mathbb{N}$.*

(1) *For any convex subgroup $H \subseteq G$, we have*

$$H + nG = \bigcap_{\substack{\alpha' \in \mathcal{S}_n \\ G_{\alpha'} \supseteq H}} (G_{\alpha'} + nG).$$

(2) *For $\alpha \in \mathcal{A}$, we have*

$$G_\alpha^{[n]} = \bigcap_{\substack{\alpha' \in \mathcal{S}_n \\ \alpha' > \alpha}} (G_{\alpha'} + nG).$$

Proof. (1) "$\subseteq$" is clear, so suppose now $a \notin H + nG$. Set $\alpha' = \mathfrak{s}_n(a)$. Then by definition $a \notin G_{\alpha'} + nG$.

(2) Again, "$\subseteq$" is clear. By applying (1) to the groups $H + nG$ appearing in the definition of $G_\alpha^{[n]}$, we obtain that $G_\alpha^{[n]}$ is the intersection of groups $G_{\alpha'} + nG$ for some $\alpha' \in \mathcal{S}_n$. Since $G_\alpha \subsetneq H \subseteq G_{\alpha'}$, these α' satisfy $\alpha' > \alpha$. $\qquad\square$

The relations $\equiv_{m,\alpha}$ and $\equiv_{m,\alpha}^{[n]}$ have a lot of similar basic properties. The following three lemmas list those which we will need; we formulate them in terms of the groups $G_\alpha + mG$ and $G_\alpha^{[n]} + mG$.

Lemma 2.5. *For $\alpha \in \mathcal{A}$ and $m, n \in \mathbb{N}$, we have*

$$G_\alpha^{[n]} + mG = G_\alpha^{[n]} + \gcd(m, n)G.$$

In particular, in L_{qe} we only need those predicates $\equiv_{m,\alpha}^{[n]}$ with $m \mid n$.

Proof. Since $nG \subseteq G_\alpha^{[n]}$, the left hand side contains $nG + mG = \gcd(m,n)G$. $\qquad\square$

Lemma 2.6. *For $k \in \mathbb{Z}$, $m,n \in \mathbb{N}$, and $\alpha \in \mathcal{A}$, we have:*

$$k(G_\alpha + mG) = kG \cap (G_\alpha + kmG)$$
$$k(G_\alpha^{[n]} + mG) = kG \cap (G_\alpha^{[kn]} + kmG)$$

Proof. Straight forward, using that the convexity of G_α implies $kG_\alpha = kG \cap G_\alpha$ and using the definition of $G_\alpha^{[n]}$. $\qquad\square$

Lemma 2.7. *Suppose that $m = m_1 \cdot m_2, n = n_1 \cdot n_2 \in \mathbb{N}$ with m_1, m_2 coprime and n_1, n_2 coprime, and suppose that $\alpha \in \mathcal{A}$. Then we have:*

$$G_\alpha + mG = (G_\alpha + m_1 G) \cap (G_\alpha + m_2 G)$$
$$G_\alpha^{[n]} + mG = (G_\alpha^{[n]} + m_1 G) \cap (G_\alpha^{[n]} + m_2 G)$$
$$G_\alpha^{[n]} + mG = (G_\alpha^{[n_1]} + mG) \cap (G_\alpha^{[n_2]} + mG)$$

Proof. The first two equations are simply the Chinese remainder theorem in the groups G/G_α and $G/G_\alpha^{[n]}$, respectively. The third one also follows directly from the Chinese remainder theorem, but since this is slightly more subtle, let us write down the details. "$\subseteq$" is clear. For "$\supseteq$", suppose that a is an element of the right hand side, i.e., there are elements $b_i \in mG$, $c_{\alpha',i} \in G_{\alpha'}$ and $d_{\alpha',i} \in G$ such that for $i = 1, 2$ and for all $\alpha' > \alpha$ we have

$$a = b_i + c_{\alpha',i} + n_i d_{\alpha',i}.$$

Find $x_1, x_2 \in \mathbb{Z}$ with $x_1 n_1 + x_2 n_2 = 1$. Then

$$a = x_1 n_1 (b_2 + c_{\alpha',2} + n_2 d_{\alpha',2}) + x_2 n_2 (b_1 + c_{\alpha',1} + n_1 d_{\alpha',1})$$
$$= \underbrace{x_1 n_1 b_2 + x_2 n_2 b_1}_{\in mG} + \underbrace{x_1 n_1 c_{\alpha',2} + x_2 n_2 c_{\alpha',1}}_{\in G_{\alpha'}} + \underbrace{n_1 n_2 (x_1 d_{\alpha',2} + x_2 d_{\alpha',1})}_{\in nG},$$

i.e., $a \in G_\alpha^{[n]} + mG$. $\qquad\square$

Let us end this section by relating the L_{syn}-maps $\mathfrak{s}_n$ and $\mathfrak{t}_n$ with the L_{qe}-predicates $\equiv_{n,\alpha}$ and $=_\alpha$.

Lemma 2.8. *For $n \in \mathbb{N}$, $a \in G$, $\alpha \in \mathcal{A}$ and $\beta \in \mathcal{S}_n \cup \mathcal{T}_n$, we have the following equivalences, where for $\overset{(1)}{\Longrightarrow}$, we additionally need $a \notin nG$,*

and for $\overset{(3)}{\Longrightarrow}$, *we additionally need* $a \neq 0$.

$$\mathfrak{s}_n(a) \geq \alpha \overset{(1)}{\Longleftrightarrow\!\!\!\!\!\Longrightarrow} a \not\equiv_{n,\alpha} 0 \qquad \mathfrak{t}_n(a) \geq \beta \overset{(3)}{\Longleftrightarrow\!\!\!\!\!\Longrightarrow} a \neq_\beta 0$$

$$\mathfrak{s}_n(a) \leq \alpha \overset{(2)}{\Longleftrightarrow} a \equiv_{n,\alpha}^{[n]} 0 \qquad \mathfrak{t}_n(a) \leq \beta \overset{(4)}{\Longleftrightarrow} a =_{\beta+} 0.$$

Proof. (1) For any convex subgroup $H \subseteq G$, we have the equivalence $G_{\mathfrak{s}_n(a)} \supseteq H \iff a \notin H + nG$, where for "$\Longrightarrow$", we additionally assume $a \notin nG$. Set $H := G_\alpha$.

(2) If $a \in nG$, then both sides are true anyway. Otherwise, (2) follows from (1) using that the right hand side is equivalent to $a \equiv_{n,\alpha'} 0$ for all $\alpha' > \alpha, \alpha' \in \mathcal{S}_n$ by Lemma 2.4 (2).

(3) If $H \subseteq G$ is a union of groups of the form G_α for $\alpha \in \mathcal{S}_n$, then we have the equivalence $G_{\mathfrak{t}_n(a)} \supseteq H \iff a \notin H$, where for "$\Longrightarrow$", we additionally assume $a \neq 0$. Set $H := G_\beta$.

(4) Again, for $a = 0$ both sides are true anyway and for $a \neq 0$, the statement follows from (3). $\qquad\square$

2.3. More dimensions of $\mathbb{F}_p$-vector spaces.

In the definition of L_{qe}, we added predicates for the dimension as $\mathbb{F}_p$-vector spaces of certain quotients of groups of the form $G_\alpha + pG$ or $G_\alpha^{[p^s]} + pG$; in particular, we required $\alpha \in \mathcal{S}_p$. The following lemma shows that this is enough to get the dimension of arbitrary quotients of two groups of this type, and for any $\alpha \in \mathcal{A}$. Moreover, we also want to consider the quotient of G by such a group. To simplify formulating the lemma, we temporarily introduce the following notation.

Notation 2.9. Set $G_\alpha^{[p^\infty]} := G_\alpha$ and $G_\infty := G$.

Note that all groups we are interested in form a long chain: for $\alpha, \alpha' \in \mathcal{A}$ with $\alpha < \alpha'$, we have

$$\cdots \subseteq G_\alpha^{[p^\infty]} + pG \subseteq \cdots \subseteq G_\alpha^{[p^2]} + pG \subseteq G_\alpha^{[p]} + pG \subseteq \ldots$$

$$\cdots \subseteq G_{\alpha'}^{[p^\infty]} + pG \subseteq \cdots\cdots \subseteq G_\infty.$$

Thus taking a quotient $(G_{\alpha_2}^{[p^{s_2}]} + pG)/(G_{\alpha_1}^{[p^{s_1}]} + pG)$ makes sense iff

$$(*) \qquad\qquad \alpha_1 < \alpha_2 \quad \vee \quad (\alpha_1 \asymp \alpha_2 \wedge s_1 \geq s_2)$$

holds.

Lemma 2.10. *Fix $p \in \mathbb{P}$, $s_1, s_2 \in \mathbb{N} \cup \{\infty\}$, and $\ell \in \mathbb{N}_0$, and fix two auxiliary sorts $\mathcal{A}_1$ and $\mathcal{A}_2$. We additionally allow $\mathcal{A}_2 = \{\infty\}$. Then the set*

$$\{(\alpha_1, \alpha_2) \in \mathcal{A}_1 \times \mathcal{A}_2 \mid (*) \text{ holds and } \dim_{\mathbb{F}_p}(G_{\alpha_2}^{[p^{s_2}]} + pG)/(G_{\alpha_1}^{[p^{s_1}]} + pG) = \ell\}$$

is $\mathcal{M}$-qf definable in L_{qe}.

Proof. Set $H_i := G_{\alpha_i}^{[p^{s_i}]} + pG$ for $i = 1, 2$. To obtain definability of the dimension of H_2/H_1 (in the above sense), it suffices to find some intermediate groups such that the dimensions of successive quotients are definable in the same sense; we will use this method to reduce to dimensions which are given by L_{qe}-predicates.

We will use Lemma 2.4 several times to show that some groups of the form $G_\alpha + pG$ or $G_\alpha^{[p]}$ are equal. By that lemma, such groups are intersections of groups $G_\beta + pG$ for some $\beta \in \mathcal{S}_p$ (note that we do not require $\alpha \in \mathcal{S}_p$), so we get equality as soon as the corresponding sets of β are equal.

Suppose first that $\alpha_1 \asymp \alpha_2$. If there is no $\beta \in \mathcal{S}_p$ with $\alpha_i \asymp \beta$, then $G_{\alpha_i} + pG = G_{\alpha_i}^{[p]}$ (by Lemma 2.4), which implies $H_1 = H_2$, since $G_{\alpha_i} + pG \subseteq H_i \subseteq G_{\alpha_i}^{[p]}$. Thus we may suppose that $\alpha_i \asymp \beta$ for some $\beta \in \mathcal{S}_p$. Moreover, we may suppose $s_1 > s_2$. If $s_1 = \infty$, then "$\dim_{\mathbb{F}_p} H_2/H_1 = \ell$" itself is a predicate of L_{qe}; otherwise compute the dimension using the chain of groups

$$H_1 \ \subseteq \ G_\beta^{[p^{s_1-1}]} + pG \ \subseteq \ G_\beta^{[p^{s_1-2}]} + pG \ \subseteq \cdots \subseteq \ H_2.$$

Now consider the case $\alpha_1 < \alpha_2$. Set $I := \{\beta \in \mathcal{S}_p \mid \alpha_1 < \beta < \alpha_2\}$. We claim that if I has cardinality bigger than $\ell+1$, then $\dim_{\mathbb{F}_p} H_2/H_1 > \ell$. Indeed, for any $\beta, \beta' \in I$ with $\beta < \beta'$, if we take $a \in G$ with $\mathfrak{s}_p(a) = \beta$, we have $a \in (G_{\beta'} + pG) \setminus (G_\beta + pG)$ so we get a strictly ascending chain of more than $\ell + 1$ groups between H_1 and H_2.

Finally, if $I = \{\beta_1, \ldots, \beta_k\}$, we use the following chain to compute the dimension:

$$H_1 \ \subseteq \ G_{\alpha_1}^{[p]} \ = \ G_{\beta_1} + pG \ \subseteq \ G_{\beta_1}^{[p]} \ = \ G_{\beta_2} + pG \ \subseteq \cdots$$
$$\cdots \subseteq \ G_{\beta_k}^{[p]} \ = \ G_{\alpha_2} + pG \ \subseteq \ H_2$$

Here, all equalities follow from Lemma 2.4, the dimension at the first and the last step have already been computed above, and the remaining dimensions are given by L_{qe}-predicates. $\square$

2.4. The predicates $x =_\bullet k_\bullet$, $x \equiv_{m,\bullet} k_\bullet$ and $D_{p^r}^{[p^s]}(x)$.

The L_{syn}-predicates $x =_\bullet k_\bullet$ and $x \equiv_{m,\bullet} k_\bullet$ were defined using quantification over all convex subgroups H of G such that G/H is discrete. The following lemma shows that this is definable.

Lemma 2.11. *If $H \subseteq G$ is any convex subgroup such that G/H is discrete, then in each of the sorts $\mathcal{S}_n, \mathcal{T}_n, n \geq 2$, there exists an α with $H = G_\alpha$. In particular:*

(1) $x =_{\bullet} k_{\bullet}$ and $x \equiv_{m,\bullet} k_{\bullet}$ are definable in L_{oag}.
(2) In any auxiliary sort, the set of α such that G/G_{α} is discrete is $\mathcal{M}$-qf definable (both, in L_{qe} and in L_{syn}).

Proof. Suppose that G/H is discrete, and choose any $a \in G$ in the preimage of the smallest positive element of G/H. Then $a \notin H + nG$ for any $n \geq 2$, but $a \in H' \subseteq H' + nG$ for any convex $H' \supsetneq H$; hence $H = G_{\mathfrak{s}_n(a)}$. Moreover, since $a \in H' \setminus H$, we also have $H = G_{\mathfrak{t}_n(a)}$.
(1) $x =_{\bullet} k_{\bullet}$ iff $\exists(\alpha \in \mathcal{S}_2)\,(\mathrm{discr}(\alpha) \wedge x =_{\alpha} k_{\alpha})$; and similarly for $\equiv_{m,\bullet}$.
(2) G/G_{α} is discrete iff $\exists(\beta \in \mathcal{S}_2)\,(\beta \asymp \alpha \wedge \mathrm{discr}(\beta))$. $\qquad\square$

The following lemma shows the connection between the L_{syn}-predicates $x =_{\bullet} k_{\bullet}$, $x \equiv_{m,\bullet} k_{\bullet}$ and $D_{p^r}^{[p^s]}$ and the corresponding L_{qe}-predicates. Each of these L_{syn}-predicates defines a union of some sets X_{α} given by the corresponding L_{qe}-predicate, where α runs through a certain auxiliary set Ξ. The point is that if x lies in this union, then α can be recovered from x by a definable function form the union to Ξ. This will allow us to define the sets X_{α} using the corresponding L_{syn}-predicate.

Lemma 2.12. *For $x \in G$ we have the following implications (1a), (2a), (3a), which in particular imply the equivalences (1b), (2b), (3b).*
(1) *For $k \in \mathbb{Z} \setminus \{0\}$ and $\alpha \in \mathcal{A}$:*

(1a) $$\mathrm{discr}(\alpha) \wedge x =_{\alpha} k_{\alpha} \Longrightarrow \alpha \asymp \mathfrak{t}_2(x)$$

(1b) $$x =_{\bullet} k_{\bullet} \iff \mathrm{discr}(\mathfrak{t}_2(x)) \wedge x =_{\mathfrak{t}_2(x)} k_{\mathfrak{t}_2(x)}$$

(2) *For $m \in \mathbb{N}, k \in \{1, \ldots, m-1\}$ and $\alpha \in \mathcal{A}$:*

(2a) $$\mathrm{discr}(\alpha) \wedge x \equiv_{m,\alpha} k_{\alpha} \Longrightarrow \alpha \asymp \mathfrak{s}_m(x)$$

(2b) $$x \equiv_{m,\bullet} k_{\bullet} \iff \mathrm{discr}(\mathfrak{s}_m(x)) \wedge x \equiv_{m,\mathfrak{s}_m(x)} k_{\mathfrak{s}_m(x)}$$

(3) *For $p \in \mathbb{P}, r, s \in \mathbb{N}$ with $s \geq r$ and $\alpha \in \mathcal{S}_{p^r}$:*

(3a) $$x \equiv_{p^r,\alpha}^{[p^s]} 0 \wedge x \not\equiv_{p^r,\alpha} 0 \Longrightarrow \alpha = \mathfrak{s}_{p^r}(x)$$

(3b) $$D_{p^r}^{[p^s]}(x) \iff x \equiv_{p^r,\mathfrak{s}_{p^r}(x)}^{[p^s]} 0 \wedge x \not\equiv_{p^r,\mathfrak{s}_{p^r}(x)} 0$$

Remark 2.13. The map $\mathfrak{t}_2$ in (1) can of course be replaced by any other map $\mathfrak{t}_p$, $p \in \mathbb{P}$.

Proof of Lemma 2.12. In (1a) and (2a), discreteness of G/G_{α} and the choice of k ensures that the left hand side implies $x \notin G_{\alpha}$ (and even $x \notin G_{\alpha} + mG$ in the case of (2a)). On the other hand, we have $x \in H$ for any convex group $H \supsetneq G_{\alpha}$. This implies the corresponding right hand side. For (3a), use $x \equiv_{p^r,\alpha}^{[p^s]} 0 \Rightarrow x \equiv_{p^r,\alpha}^{[p^r]} 0$ and Lemma 2.8.

In (Xb), x satisfies the left hand side if and only there is an α (in $\mathcal{T}_2$, $\mathcal{S}_m$ or $\mathcal{S}_{p^r}$, respectively) such that x satisfies the left hand side of (Xa).

The right hand side of (Xa) says how this α can be obtained from x. Plugging this in yields the right hand side of (Xb). $\qquad\qquad\square$

2.5. **Translation between L_{syn} and L_{qe}.** When introducing the language L_{syn}, we claimed that it is strong enough to express L_{qe} without $\mathcal{M}$-quantifiers. On the other hand, we want to deduce quantifier elimination in L_{qe} from quantifier elimination in L_{syn}, hence we also need (a version of) the other direction. This is what we prove in this section. At the end of the section, the translation $L_{\mathrm{syn}} \rightsquigarrow L_{\mathrm{qe}}$ will be applied to deduce Theorem 1.8 from Theorem 1.13.

Proposition 2.14. *Any L_{qe}-predicate can be expressed in L_{syn} without $\mathcal{M}$-quantifiers.*

Remark 2.15. Since any function symbol in L_{qe} is also contained in L_{syn}, this implies that any $\mathcal{M}$-qf L_{qe}-formula is equivalent to an $\mathcal{M}$-qf L_{syn}-formula.

Proof of Proposition 2.14. The predicates of $L_{\mathrm{qe}} \backslash L_{\mathrm{syn}}$ are the following:

- $x \diamond_\eta y + k_\eta$ where $\diamond \in \{=, <, \equiv_m\}$, $k \in \mathbb{Z}$, $m \in \mathbb{N}$, and where η may be from any of the sorts of $\mathcal{A}$;
- $x \equiv_{m,\eta}^{[n]} y$ for $m, n \in \mathbb{N}$ and where η is from one of the sorts $\mathcal{S}_p$.

Concerning $x \diamond_\eta y + k_\eta$, if $k \neq 0$, then we may assume that G/G_η is discrete, since otherwise by definition $k_\eta = 0$. (Recall that this discreteness is definable on any auxiliary sort by Lemma 2.11.)

We now translate all these predicates into L_{syn}, starting with the easier ones so that we can use them for the more difficult ones.

First consider $x \equiv_{m,\eta} y$ (for η from any $\mathcal{A}$-sort). By Lemma 2.8 (1), this is equivalent to $\mathfrak{s}_m(x - y) < \eta \vee x \equiv_m y$, which is equivalent to $\bigwedge_{p^r || m} \mathfrak{s}_{p^r}(x - y) < \eta \vee x \equiv_m y$ by Lemma 2.2.

Next consider $x =_\eta y$. If $\eta \in \mathcal{S}_p \cup \mathcal{T}_p$, then by Lemma 2.8 (3) this is equivalent to $\mathfrak{t}_p(x - y) < \eta \vee x = y$. If $\eta \in \mathcal{T}_p^+$, then it is equivalent to $\forall (\theta \in \mathcal{S}_p)\, (\theta \geq \eta \to x =_\theta y)$.

Now consider $x =_\eta y + k_\eta$ for $k \neq 0$. (Recall that we assume now that G/G_η is discrete.) Then Lemma 2.12 (1a) implies $\eta = \mathfrak{t}_2(x - y)$, and under this assumption, Lemma 2.12 (1b) implies that $x =_\eta y + k_\eta$ is equivalent to $x - y =_\bullet k_\bullet$. Thus (under the assumption $\mathrm{discr}(\eta)$):

$$x =_\eta y + k_\eta \iff \eta = \mathfrak{t}_2(x - y) \wedge x - y =_\bullet k_\bullet.$$

Exactly the same argument yields, for $m \in \mathbb{N}$ and $k \in \{1, \ldots, m-1\}$ (which we may assume):

$$x \equiv_{m,\eta} y + k_\eta \iff \eta = \mathfrak{s}_m(x - y) \wedge x - y \equiv_{m,\bullet} k_\bullet.$$

Concerning $x \equiv^{[n]}_{m,\eta} y$, we may assume that m and n are prime powers by Lemma 2.7, and we may assume $m \mid n$ by Lemma 2.5; so $m = p^r$ and $n = p^s$ for some $p \in \mathbb{P}$ and $s \geq r$. Moreover, it suffices to define $x \equiv^{[p^s]}_{p^r,\eta} y \wedge \neg x \equiv_{p^r,\eta} y$; this again works in the same way as before with Lemma 2.12, yielding:

$$x \equiv^{[p^s]}_{p^r,\eta} y \iff x \equiv_{p^r,\eta} y \vee (\eta = \mathfrak{s}_{p^r}(x - y) \wedge D^{[p^s]}_{p^r}(x - y)).$$

Finally, consider $x <_\eta y + k_\eta$. If $k = 0$, then this is equivalent to $x < y \wedge x \neq_\eta y$. If k is positive, then we take the disjunction of this with $x =_\eta y + i_\eta$ for $0 \leq i < k$; if k is negative, then we take the conjunction of this with $x \neq_\eta y + i_\eta$ for $k \leq i < 0$. $\qquad\square$

Proposition 2.16. *Every quantifier free L_{syn}-formula is equivalent to an L_{qe}-formula in family union form.*

Proof. Let $\phi(\bar{x}, \bar{\eta})$ be a given quantifier free L_{syn}-formula; we have to get rid of the following kind of atoms:

(1) $t_1 \diamond t_2$ where $\diamond \in \{<, >, \equiv_m\}$ and t_1, t_2 are main sort terms (and $m \in \mathbb{N}$);

(2) $t =_\bullet k_\bullet$, $t \equiv_{m,\bullet} k_\bullet$ and $D^{[p^s]}_{p^r}(t)$, where t is a main sort term (and $p \in \mathbb{P}$, $m, r, s \in \mathbb{N}$);

(3) atoms involving $\mathfrak{s}_{p^r}(t)$ or $\mathfrak{t}_p(t)$, where t is a main sort term (and $p \in \mathbb{P}, r \in \mathbb{N}$).

An atom $t_1 \diamond t_2$ of type (1) can be replaced by $t_1 \diamond_{\mathfrak{s}_2(0)} t_2$. To get rid of the atoms of type (2), apply Lemma 2.12 (1b), (2b), (3b). It remains to get rid of the functions $\mathfrak{s}_m$ and $\mathfrak{t}_p$ (for $m \in \mathbb{N}$, $p \in \mathbb{P}$) (including the newly introduced ones) and bring the formula into family union form.

Let $\tau_i(\bar{x})$ be the terms of ϕ which are of the form $\mathfrak{s}_m(t(\bar{x}))$ or $\mathfrak{t}_p(t(\bar{x}))$, where $m \in \mathbb{N}$, $p \in \mathbb{P}$, and where $t(\bar{x})$ is a main sort term. We replace ϕ by the equivalent formula

$$\exists \bar{\theta} \left(\left(\bigwedge_i \tau_i(\bar{x}) = \theta_i \right) \wedge \phi[\tfrac{\theta_i}{\tau_i(\bar{x})}]_i \right).$$

(Here, the notation $\phi[\tfrac{r}{s}]$ means: the formula obtained from ϕ by replacing all occurrences of s by r.) The atoms $\tau_i(\bar{x}) = \theta_i$ can be expressed in L_{qe} using Lemma 2.8: $\mathfrak{s}_m(t(\bar{x})) = \theta$ is equivalent to

$$(t(\bar{x}) \not\equiv_{m,\theta} 0 \wedge t(\bar{x}) \equiv^{[m]}_{m,\theta} 0) \vee$$
$$(\theta \text{ is the minimal element of } \mathcal{S}_m \wedge t(\bar{x}) \equiv_{m,\theta} 0)$$

(the second line treats the case $t(\bar{x}) \equiv_m 0$), and $\mathfrak{t}_p(t(\bar{x})) = \theta$ is equivalent to

$$(t(\bar{x}) \neq_\theta 0 \wedge t(\bar{x}) =_{\theta+} 0) \vee$$
$$(\theta \text{ is the minimal element of } \mathcal{T}_p \wedge t(\bar{x}) = 0),$$

where $t(\bar{x}) =_{\theta+} 0$ can be written in family union form as

$$\exists(\theta' \in \mathcal{T}_p^+)\,(\theta' = \theta+ \wedge t(\bar{x}) =_{\theta'} 0).$$

(Here, we use $\mathcal{M}$-qf definability of $\theta \mapsto \theta+$).

Now our formula $\phi(\bar{x}, \bar{\eta})$ is purely in the language L_{qe} and it is of the form $\exists\bar{\theta}\,\psi(\bar{x}, \bar{\eta}, \bar{\theta})$, where $\bar{\theta}$ is auxiliary and ψ is a boolean combination of quantifier free parts and of parts living purely in $\mathcal{A}$. Moreover, by the way in which the quantifier $\exists\bar{\theta}$ has been introduced, $\psi(\bar{x}, \bar{\eta}, \bar{\alpha})$ and $\psi(\bar{x}, \bar{\eta}, \bar{\alpha}')$ are inconsistent for any $\bar{\alpha}, \bar{\alpha}' \in \mathcal{A}$, $\bar{\alpha} \neq \bar{\alpha}'$. Thus to turn $\phi(\bar{x}, \bar{\eta})$ into family union form, it remains to bring ψ into a disjunctive normal form where the conjunctive clauses are pairwise inconsistent, and then pull the disjunction to the outside (here, we treat the $\mathcal{A}$-parts of ψ with quantifiers as atoms). This kind of disjunctive normal form can be obtained by using conjunctive clauses each of which contains all atoms occurring in ψ, either positively or negatively.   □

Now it is easy to deduce L_{qe} quantifier elimination from L_{syn} quantifier elimination:

Proof of Theorem 1.8 from Theorem 1.13. Any L_{qe}-formula is equivalent to an L_{syn}-formula. Using Theorem 1.13, we can turn this into an L_{syn}-formula in family union form

$$\phi(\bar{x}, \bar{\eta}) = \bigvee_{i=1}^{k} \exists\bar{\theta}\,\big(\xi_i(\bar{\eta}, \bar{\theta}) \wedge \psi_i(\bar{x}, \bar{\theta})\big).$$

Since L_{syn} and L_{qe} agree on the auxiliary sorts, the formulas ξ_i are also L_{qe}-formulas. By Proposition 2.16, we may replace each ψ_i by an L_{qe}-formula in family union form. By pulling the quantifiers and disjunctions of these ψ_i to the outside, we obtain a formula which is in family union form as a whole.   □

3. The main proofs

3.1. **Partial quantifier elimination in general.** In this section, we prove Proposition 1.11 which gives a general method to eliminate main sort quantifiers when the only connection between the main sorts and the auxiliary sorts are functions from $\mathcal{M}$ and $\mathcal{A}$. The proof goes in two steps; we formulate the first one as a separate lemma.

Lemma 3.1. *Let L be a language, let $\mathcal{M} \,\dot\cup\, \mathcal{A}$ be a partition of the sorts of L, and suppose that the only symbols in L connecting $\mathcal{M}$ and $\mathcal{A}$ are functions from (products of) $\mathcal{M}$ sorts to $\mathcal{A}$ sorts. Then any formula without $\mathcal{M}$-quantifiers is equivalent to a formula in family union form (in any theory).*

Proof. Let ϕ be an $\mathcal{M}$-qf formula. We do an induction over the number of occurrences of main variables in ϕ. If no main variable appears in ϕ, there is nothing to do. Otherwise, choose a specific occurrence of a main variable x in ϕ. We distinguish the following two cases:

(1) The atom a containing x is a relation on $\mathcal{M}$ (applied to some terms living completely in $\mathcal{M}$).

(2) x appears inside a term t with range in $\mathcal{A}$.

In case (1), every other variable appearing in the atom a is also a main sort variable, so a does not depend on any of the quantified variables of ϕ, and we can "do a case distinction on a": ϕ is equivalent to

$$(a \wedge \phi[\tfrac{\top}{a}]) \vee (\neg a \wedge \phi[\tfrac{\bot}{a}]).$$

(Here, the notation $\phi[\tfrac{r}{s}]$ means: the formula obtained from ϕ by replacing all occurrences of s by r, $\top$ means true and $\bot$ means false.) Apply the induction hypothesis to $\phi[\tfrac{\top}{a}]$ and $\phi[\tfrac{\bot}{a}]$. After pulling the "$a\wedge$" and "$\neg a \wedge$" inside, the result is in family union form.

In case (2), consider the smallest subterm t' of t containing x whose range lies in $\mathcal{A}$. Then the outermost function of t' is a function from a product of some $\mathcal{M}$-sorts to $\mathcal{A}$, so t' depends only on $\mathcal{M}$-variables and in particular not on quantified variables. Thus ϕ is equivalent to

$$\exists \xi \, (t' = \xi \wedge \phi[\tfrac{\xi}{t'}]).$$

Applying induction to $\phi[\tfrac{\xi}{t'}]$ yields a formula in family union form. $\square$

Note that this lemma in particular implies that the negation of a formula in family union form can again be brought into family union form.

Now let us get to the main proof of this section:

Proof of Proposition 1.11. Let ϕ be a formula whose $\mathcal{M}$-quantifiers we want to eliminate. We use induction over the structure of ϕ, i.e., we suppose that the subformulas are already in family union form. By Lemma 3.1, it suffices to bring ϕ into a form without $\mathcal{M}$-quantifiers.

If ϕ is an atom, then there is nothing to do, and neither if it is of the form $\neg \psi$ or $\psi_1 \wedge \psi_2$, so suppose $\phi = \exists x \, \psi(x)$, where x is a main sort

variable and $\psi(x)$ is in family union form, i.e.:

$$\phi(\bar{y}, \bar{\eta}) = \exists x \bigvee_{i=1}^{k} \exists \bar{\theta} \left(\xi_i(\bar{\eta}, \bar{\theta}) \wedge \psi_i(x, \bar{y}, \bar{\theta}) \right)$$

Rewrite this as

$$\bigvee_{i=1}^{k} \exists \bar{\theta} \left(\xi_i(\bar{\eta}, \bar{\theta}) \wedge \exists x \, \psi_i(x, \bar{y}, \bar{\theta}) \right).$$

Since $\psi_i(x, \bar{y}, \bar{\theta})$ is quantifier free, the hypothesis of the proposition applies to $\exists x \, \psi_i(x, \bar{y}, \bar{\theta})$, and we get a formula without main sort quantifiers. $\qquad\square$

3.2. **Removing the quantifier in $X + G'$.** At one point in the main proof of quantifier elimination, we will have a subgroup $G' \subseteq G$ and a set $X \subseteq G$ defined by a quantifier free formula of a particular form and we will need to be able to define the set $X + G'$ without quantifiers. This will be possible using the following two lemmas which have nothing to do with model theory.

Lemma 3.2. *Suppose we have an abelian group G, a subgroup $G' \subseteq G$ and a subset $X \subseteq G$ of the form*

$$X = (H_0 + a_0) \setminus \bigcup_{i=1}^{\nu} (H_i + a_i)$$

where H_i are subgroups of G, $a_i \in G$, and where $H_i + a_i \subseteq H_0 + a_0$ for $i \in \{1, \dots, \nu\}$ and $H_i + a_i \cap H_j + a_j = \emptyset$ for $i, j \in \{1, \dots, \nu\}, i \neq j$. Then for $x \in G$ we have $x \in X' := X + G'$ if and only if

(1) $$x - a_0 \in H_0 + G' \quad \text{and}$$

(2) $$\sum_{\{1 \leq i \leq \nu \mid x - a_i \in H_i + G'\}} ((H_0 \cap G') : (H_i \cap G'))^{-1} < 1.$$

(Here, we use the convention $\infty^{-1} = 0$.)

Proof. The condition $b \in X'$ is equivalent to $X \cap (b + G') \neq \emptyset$. Write

$$X \cap (b + G') = C_0 \setminus \bigcup_{i=1}^{\nu} C_i,$$

with $C_i := (a_i + H_i) \cap (b + G')$. Then C_i is non-empty if and only if $b - a_i \in H_i \cap G'$, and if it is non-empty, then it is of the form $c_i + H_i \cap G'$.

Non-emptiness of C_0 is just condition (1) on b in the lemma, so suppose now that C_0 indeed is non-empty. The question is now whether the union $\bigcup_{i=1}^{\nu} C_i$ (which is disjoint) contains all of C_0. The sum in condition (2) goes exactly over those $i \geq 1$ for which C_i is non-empty,

and the summand is the proportion of C_i in C_0. Hence $\bigcup_{i=1}^{\nu} C_i = C_0$ if and only if the sum is 1. (To make this more formal, count elements in C_0/D, where D is the intersection of all those C_i which have finite index in C_0.) $\qquad\square$

The next lemma will be helpful to make condition (2) from the previous lemma definable.

Lemma 3.3. *Suppose that $n \in \mathbb{N}$, $n \geq 2$ and that $q_1, \ldots, q_{\nu}$ are powers of n. Then there exists an $N \in \mathbb{N}$ depending only on n and ν such that*

$$\sum_{i=1,\ldots,\nu} q_i^{-1} \geq 1 \quad \Longleftrightarrow \quad \sum_{\substack{i=1,\ldots,\nu \\ q_i < n^N}} q_i^{-1} \geq 1.$$

Proof. Choose N such that $\nu < N \cdot (n-1) + 1$. Without loss, $q_1 \leq \cdots \leq q_{\nu}$. Set $s_k := \sum_{i=1}^{k} q_i^{-1}$ and let d_k be the digit sum of s_k in base n. Inductively, one proves $d_k \leq k$. If the claim of the lemma is false, then there exists an $\ell \leq \nu$ with $q_\ell \geq n^N$ such that $s_{\ell-1} < 1 \leq s_\ell$. This implies $d_{\ell-1} \geq N \cdot (n-1)$, contradicting $d_{\ell-1} \leq \ell-1 \leq \nu-1 < N \cdot (n-1)$. $\quad\square$

3.3. **Actually eliminating the quantifiers.**

Proof of Theorem 1.13. As announced, we prove Theorem 1.13 using Proposition 1.11, i.e., we have to show that if $\phi(x, \bar{y}, \bar{\eta})$ is a quantifier free L_{syn}-formula, then $\exists x\, \phi(x, \bar{y}, \bar{\eta})$ is equivalent to an $\mathcal{M}$-qf L_{syn}-formula. Since the language L_{qe} is more intuitive, we start by translating ϕ into an L_{qe}-formula using Proposition 2.16. The result is in family union form, i.e., we have to eliminate "$\exists x$" from a formula of the form

$$\exists x \bigvee_{i=1}^{k} \exists \bar{\theta} \left(\xi_i(\bar{\eta}, \bar{\theta}) \wedge \psi_i(x, \bar{y}, \bar{\theta}) \right).$$

By pulling this quantifier inside, it suffices to eliminate the quantifier of $\exists x\, \psi_i(x, \bar{y}, \bar{\theta})$. In other words, we now need to eliminate the quantifier of $\exists x\, \phi(x, \bar{y}, \bar{\eta})$ when $\phi(x, \bar{y}, \bar{\eta})$ of the form

$$(*) \qquad\qquad \phi(x, \bar{y}, \bar{\eta}) = \bigwedge_{i=1}^{k} r_i x \, (\diamond_i)_{\eta_i} \, y_i + k_{\eta_i}$$

with $r_i \in \mathbb{N}$, $\diamond_i \in \{=, \neq, <, >, \leq, \geq, \equiv_m, \not\equiv_m, \equiv_m^{[n]}, \not\equiv_m^{[n]}\}$.

We will show that $\exists x\, \phi(x, \bar{y}, \bar{\eta})$ is equivalent to an $\mathcal{M}$-qf formula in the language $L_{\mathrm{syn}} \cup L_{\mathrm{qe}}$; this is enough, since afterwards, we can apply Proposition 2.14 to translate the L_{qe}-predicates into L_{syn}.

To simplify the exposition, let us choose parameters $\bar{b} \in \mathcal{M}$, $\bar{\alpha} \in \mathcal{A}$ and consider $\phi(x, \bar{b}, \bar{\alpha})$; we will denote this by $\phi(x)$ for short. Our

strategy is to successively simplify $\phi(x)$; of course, the whole point is that this is done in a way depending definably on the parameters $\bar{b}$ and $\bar{\alpha}$.

If $x \diamond_\alpha b + k_\alpha$ is a literal of ϕ, we will write b_* for a representative in G of $b + k_\alpha$, so that $x \diamond_\alpha b + k_\alpha$ is equivalent to $x \diamond_\alpha b_*$; if $k = 0$ or G/G_α is dense, we set $b_* := b$. We will sometimes use $x \diamond_\alpha b_*$ as a short hand notation for $x \diamond_\alpha b + k_\alpha$. Of course, we are not allowed to use b_* in the resulting quantifier free formula. However, if we have an element $\alpha' \geq \alpha$, then a condition of the form $t \diamond'_{\alpha'} b_*$ can easily be expressed using b and k instead of b_*:

$$t \diamond'_{\alpha'} b_* \quad \Longleftrightarrow \quad (\alpha' = \alpha \wedge t \diamond'_\alpha b + k_\alpha) \vee (\alpha' > \alpha \wedge t \diamond'_{\alpha'} b);$$

we will use this without further mentioning.

Now let us get to work. First we get rid of the factors r_i in (*). To this end, note that in each literal, multiplying both sides by any non-zero integer r doesn't change the set defined by that literal if additionally we do the following:

- in literals with $\equiv_m$, $\not\equiv_m$, $\equiv_m^{[n]}$, $\not\equiv_m^{[n]}$, we also multiply m (and n) by r (this uses Lemma 2.6);
- we turn inequalities around if $r < 0$.

In this way, we can make all r_i equal to one single r. After that, we replace rx by a new variable x' and replace $\exists x$ by $\exists (x' \in rG)$.

The remainder of the proof will consist of two big parts: in the first one, we get rid of the inequalities ($\neq, <, >, \leq, \geq$); in the second one, we treat the congruence conditions ($\equiv_m, \not\equiv_m, \equiv_m^{[n]}, \not\equiv_m^{[n]}$).

Part 1: treating inequalities

Our goal in this part is to reduce the quantifier elimination problem to formulas $\phi(x)$ of the form

$$(**) \qquad\qquad \phi'(x) \quad \text{or} \quad x =_\delta b_* \wedge \phi'(x),$$

where the atoms of $\phi'(x)$ use only $\equiv_m$ and $\equiv_m^{[n]}$.

We start by replacing literals of the form $x =_\alpha b_*$ and $x \neq_\alpha b_*$ by $x \geq_\alpha b_* \wedge x \leq_\alpha b_*$ and $x >_\alpha b_* \vee x <_\alpha b_*$, respectively. In the second case, we treat each disjunct separately. (Replacing $x =_\alpha b_*$ by inequalities might seem strange at first sight, since later, we want to get back to equalities. However, recall that after all, $x =_\alpha b_*$ defines an interval.)

Next, reduce to the case where $\phi(x)$ contains at most one lower and one upper bound: if $\phi = \phi'' \wedge \psi_1 \wedge \psi_2$, where ψ_i are two bounds on the same side, then $\exists x \, \phi$ is equivalent to $(\exists x \, (\phi'' \wedge \psi_1)) \wedge (\exists x \, (\phi'' \wedge \psi_2))$.

Thus $\phi(x)$ is now of the form

$$(\text{***}) \qquad\qquad \phi(x) = \quad c_* \lhd_\alpha x \lhd'_{\alpha'} c'_* \wedge \phi'(x),$$

where $\lhd, \lhd' \in \{\leq, <, \text{no condition}\}$ and where the atoms of $\phi'(x)$ use only $\equiv_m$ or $\equiv_m^{[n]}$. Such an atom defines a union of cosets of mG, hence if we let m_0 be the least common multiple of all occurring m, then $\phi'(G)$ is a union of cosets of $m_0 G$. We fix this m_0 for the remainder of the proof.

If $\phi(x)$ has no bounds, then it is already of the form (**). If $\phi(x)$ has only one bound, say, a lower one, then removing that bound does not change the truth of $\exists x \, \phi(x)$. Indeed, if an element $a \in G$ satisfies $\phi'(x)$ but does not satisfy $c_* \lhd_\alpha x$, all elements of $a + m_0 G$ satisfy $\phi'(x)$ and in that set, we can find one which also satisfies the bound. Hence for the remainder of part 1, we assume that $\phi(x)$ has two bounds.

If $\alpha \geq \alpha'$ in (***), we may suppose that $c_* \lhd_\alpha c'_*$, since otherwise $\phi(x)$ defines the empty set. Similarly, if $\alpha' \geq \alpha$, we may suppose that $c_* \lhd'_{\alpha'} c'_*$.

Let γ be an auxiliary element satisfying $\gamma \asymp \max\{\alpha, \alpha', \mathsf{t}_{m_0}(c_* - c'_*)\}$. Recall that c_* is a representative of $c + k_\alpha$ (for some $c \in G$ and $k \in \mathbb{Z}$), but since $\mathsf{t}_{m_0}(c - c_*) \leq \alpha$, γ does not depend on the choice of c_* (and similarly for c'_*). By Lemma 2.3, γ is definable. (Formally, it is an element of one of finitely many auxiliary sorts; we do a case distinction on the sort.)

Suppose that $\alpha < \gamma$. We claim that then, weakening the lower bound from $c_* \lhd_\alpha x$ to $c_* \leq_\gamma x$ does not change the truth value of the formula $\exists x \, \phi(x)$. In other words, we claim that if there exists an $a \in G$ with $a =_\gamma c_* \wedge a \lhd'_{\alpha'} c'_* \wedge \phi'(a)$, then we can find an a' satisfying $\phi(x)$. If $c_* \lhd_\alpha a$, there is nothing to do. Otherwise, we can choose an element $a_0 \in m_0 G_\gamma$ such that $c_* \lhd_\alpha a + a_0 =: a'$. By construction, a' satisfies $\phi'(x)$ and the lower bound. Concerning the upper bound: if $\gamma = \alpha'$, then $a' =_\gamma c_* \lhd'_{\alpha'} c'_*$ implies $a' \lhd'_{\alpha'} c'_*$. If, on the other hand, $\gamma > \alpha'$, then by definition of γ we have $\gamma = \mathsf{t}_{m_0}(c_* - c'_*)$, hence $c_* \neq_\gamma c'_*$, and hence $a' =_\gamma c_* <_\gamma c'_*$, which again implies $a' \lhd'_{\alpha'} c'_*$.

We do the same with the upper bound and thus get a formula of the form

$$c_* \lhd_\gamma x \lhd'_\gamma c'_* \wedge \phi'(x),$$

where $\gamma \geq \mathsf{t}_{m_0}(c_* - c'_*)$.

Now we distinguish two cases, depending on whether $c'_* - c_* >_\gamma (m_0 + 1)_\gamma$ or not. (Recall that if G/G_γ is dense, then by definition this is equivalent to $c'_* - c_* >_\gamma 0$.)

Suppose first that the condition is false. If G/G_γ is dense, then this implies $c_* =_\gamma c'_*$, so $\phi(x)$ can only be consistent if both inequalities are non-strict, and in that case, it is equivalent to $x =_\gamma c_* \wedge \phi'(x)$, which is of the form (**). If G/G_γ is discrete, then $c'_* - c_* =_\gamma \ell_\gamma$ for some $\ell \le m_0 + 1$. Thus $\phi(x)$ is equivalent to the disjunction of finitely many formulas of the form

$$x =_\gamma c_* + i_\gamma \wedge \phi'(x).$$

More precisely, i runs from 0 or 1 (depending on $\lhd$) to $\ell - 1$ or ℓ (depending on $\lhd'$).

Now suppose that $c'_* - c_* >_\gamma (m_0 + 1)_\gamma$. Then there exists an element $d \in G$ satisfying $0 <_\gamma (m_0 + 1)d <_\gamma c'_* - c_*$. (If G/G_γ is discrete, then choose for d any representative of 1_γ.) Using this, we will show that $\exists x \, \phi(x)$ is equivalent to

$$\exists x \, (x =_\delta c_* \wedge \phi'(x)),$$

where $\delta := (\mathfrak{t}_{m_0}(c'_* - c_*))+$. (Again, δ is definable by Lemma 2.3.)

It is clear that $\phi(x)$ implies $x =_\delta c_*$, since $c'_* =_\delta c_*$, so it remains to show that if there exists an $a \in c_* + G_\delta$ satisfying $\phi'(x)$, then there exists an $a' \in G$ which additionally lies between the bounds.

The inequality $\mathfrak{s}_{m_0}(a - c_*) \le \mathfrak{t}_{m_0}(a - c_*) \le \mathfrak{t}_{m_0}(c_* - c'_*) \le \gamma$ means that for any convex subgroup $H \subseteq G$ strictly containing G_γ, we have $a - c_* \in H + m_0 G$. In particular, since $d \ne_\gamma 0$,

$$a - c_* \in \langle d \rangle^{\mathrm{conv}} + m_0 G = [d, (m_0 + 1)d] + m_0 G.$$

Choose $a_0 \in (a - c_* + m_0 G) \cap [d, (m_0 + 1)d]$ and set $a' := a_0 + c_*$. Then a' satisfies $\phi'(x)$ since it differs from a by an element of $m_0 G$, and $0 <_\gamma d \le a_0 \le (m_0 + 1)d <_\gamma c'_* - c_*$ implies that a' also satisfies the bounds.

Part 2: treating congruences

Our formula $\phi(x)$ is now of the form

$$\phi'(x) \quad \text{or} \quad x =_\gamma c_* \wedge \phi'(x),$$

where the atoms of $\phi'(x)$ are of the form $x \equiv_{m,\alpha} b_*$ or $x \equiv_{m,\alpha}^{[n]} b_*$. Using Lemma 2.7, we can suppose that each m and each n is a power of a prime, and using Lemma 2.5, we get rid of all those atoms $x \equiv_{m,\alpha}^{[n]} b$ where m and n are powers of different primes. By the Chinese remainder theorem, we can eliminate the quantifier separately for each of the subformulas of ϕ' corresponding to the different primes. In other words, we may assume that all atoms of ϕ' are of the form $x \equiv_{p^r,\alpha} b_*$ or $x \equiv_{p^r,\alpha}^{[p^s]} b_*$ for one single prime p which we fix for the remainder of the proof. Moreover, in $\equiv_{p^r,\alpha}^{[p^s]}$ we may assume $s \ge r$ (again by Lemma 2.5).

From now on, we also fix r to be the maximal exponent of p appearing in the atoms in the above way (both, in $\equiv_{p^r,\alpha}$ and in $\equiv_{p^r,\alpha}^{[p^s]}$); in particular, $\phi'(G)$ consists of entire cosets of $p^r G$.

In general, if $\phi(x) = \phi_0(x) \wedge \phi_1(x)$ and $H \subseteq G$ is any subgroup such that $\phi_1(G)$ consists of entire cosets of H, then replacing ϕ_0 by a formula defining $\phi_0(G) + H$ does not change the truth of $\exists x\, \phi(x)$; we will apply this enlargement argument several times. Since $\phi'(G)$ is a union of cosets of $p^r G$, we already can replace $x =_\gamma c_*$ by $x \equiv_{p^r,\gamma} c_*$, i.e., without loss there is no literal $x =_\gamma c_*$.

Now we prove quantifier elimination by induction on r. If $r = 0$, then $\exists x\, \phi(x)$ is equivalent to $\phi(0)$. For the induction step, suppose $r > 0$ and write $\phi = \phi_0 \wedge \phi_1$, where ϕ_0 contains the atoms $x \equiv_{m,\alpha} b_*$, $x \equiv_{m,\alpha}^{[n]} b_*$ with $m = p^r$ and ϕ_1 contains the atoms with $m \leq p^{r-1}$. By the enlargement argument, we are done with the induction step if we can show the following:

(a) the set $\phi_0(G) + p^{r-1}G$ is definable by a formula ϕ_0' using only atoms of the form $x \equiv_{p^{r-1},\alpha} b_*$, $x \equiv_{p^{r-1},\alpha}^{[s]} b_*$, with r as given and $s \geq r - 1$ arbitrary;

(b) ϕ_0' depends on the parameters of ϕ_0 in an $\mathcal{M}$-qf definable way.

The atoms $x \equiv_{p^r,\alpha} b_*$, $x \equiv_{p^r,\alpha}^{[p^s]} b_*$ of ϕ_0 define cosets of groups, and these groups are totally ordered by inclusion:

$$\cdots \subseteq G_\alpha + p^r G \subseteq \cdots \subseteq G_\alpha^{[p^{s+1}]} + p^r G \subseteq G_\alpha^{[p^s]} + p^r G \subseteq \cdots \subseteq G_{\alpha'} + p^r G \subseteq \ldots$$

for all $\alpha < \alpha'$ and all $s \geq r$. In particular, any two such cosets $H + b_*$, $H' + b_*'$ are either disjoint or contained in one another. Moreover, whether $H + b_* \subseteq H' + b_*'$ or not is definable. Using this, we can simplify ϕ_0 such that it has at most one positive literal, and all negative literals exclude pairwise disjoint sets. Now $\phi(G)$ satisfies the prerequisites of Lemma 3.2: in that lemma, let $H_0 + a_0$ be the set defined by the positive literal of ϕ_0 (or $H_0 = G, a_0 = 0$ if there is no positive literal), let $H_i + a_i$ be the sets excluded by the negative literals, and set $G' := p^{r-1}G$. (The a_i are the representatives denoted by b_* before.) To get our desired formula defining $X' = \phi_0(G) + p^{r-1}G$, it remains to verify that conditions (1) and (2) of Lemma 3.2 are $\mathcal{M}$-qf definable, where x only appears in atoms as in (a).

For each $i \in \{0, \ldots, \nu\}$ we have

$$H_i = G \quad \text{or} \quad H_i = G_\alpha + p^r G \quad \text{or} \quad H_i = G_\alpha^{[p^s]} + p^r G,$$

so the condition $x - a_i \in H_i + p^{r-1}G$ is definable by

$$x = x \quad \text{or} \quad x \equiv_{p^{r-1},\alpha} a_i \quad \text{or} \quad x \equiv_{p^{r-1},\alpha}^{[p^s]} a_i.$$

This settles definability of (1), and it allows us to do a case distinction which fixes the set the sum (2) runs over. Let I be that set and set $q_i := \big((H_0 \cap p^{r-1}G) : (H_i \cap p^{r-1}G)\big)$ for $i \in I$. By Lemma 2.6, for $i \in I \cup \{0\}$ we can write $H_i \cap p^{r-1}G$ as $p^{r-1}H_i'$ with

$$H_i' = G \quad \text{or} \quad H_i' = G_\alpha + pG \quad \text{or} \quad H_i' = G_\alpha^{[p^{s-r+1}]} + pG,$$

so $q_i = (H_0' : H_i')$ is the cardinality of a quotient treated by Lemma 2.10. Thus each q_i is either infinite or a power of p, and the conditions $q_i = p^\ell$ (for $\ell \in \mathbb{N}_0$) are $\mathcal{M}$-qf definable. By Lemma 3.3, there exists a bound N such that $\sum_{i \in I} q_i < 1$ iff $\sum_{i \in I, q_i < p^N} q_i < 1$. The latter is equivalent to a finite boolean combination of conditions of the form $q_i = p^\ell$ (for $i \in I$ and $\ell < N$). Hence, condition (2) is definable, too, and we are done. $\square$

4. Examples

In this section, we give some examples which should help the reader understanding the languages which we define. These examples show that large parts of the languages is indeed necessary. More detailed examples explicitly concerning L_{qe} are given in [6]; similar motivating examples, but presented from a different point of view are given in [10].

4.1. **Concrete examples illustrating the sort $\mathcal{S}_n$.** Set $G = \mathbb{Z} \oplus \mathbb{Z}$ with lexicographical order. We determine the sort $\mathcal{S}_n$ for $n \geq 2$. For this, we have to go through all elements $a \in G \setminus nG$ and find the largest convex subgroups $H = G_{\mathfrak{s}_n(a)} \subseteq G$ such that $H + nG$ does not contain a. Equivalently, H is the largest convex subgroup which is disjoint from $a + nG$.

Obviousely, H only depends on the class of a modulo nG. If $a = (0, z)$ for $z \notin n\mathbb{Z}$, then we have $H = \{(0,0)\} =: G_0$; if $a = (z, z')$ for $z \notin n\mathbb{Z}$ and $z' \in \mathbb{Z}$ arbitrary, then $H = \{0\} \times \mathbb{Z} =: G_1$. Thus $\mathcal{S}_n$ consists of two elements which correspond to the groups G_0 and G_1. (For $a \in nG$, by definition we also have $G_{\mathfrak{s}_n(a)} = G_0$.)

In this example, all sorts $\mathcal{S}_n$ are the same. Now consider the group $G = \mathbb{Z}[\frac{1}{5}] \oplus \mathbb{Z}$ instead. The sorts $\mathcal{S}_n$ for $n \neq 5^r$ are the same as before; however, the sort $\mathcal{S}_{5^r}$ now consists of a single element, since modulo $5^r G$, any element of G is equivalent to an element of the form $(0, z)$.

In these examples, the sorts $\mathcal{T}_n$ and $\mathcal{T}_n^+$ do not yield any new nontrivial convex subgroups of G: $G_{\mathfrak{t}_n(a)}$ is G_0 if $a \in G_1$ and G_1 otherwise, and $G_{\mathfrak{t}_n(a)+}$ is G_1 if $a \in G_1$ and G otherwise. To get interesting new convex subgroups, we have to consider infinite lexicographical products.

4.2. **Infinite lexicographical products illustrating $\mathcal{T}_n$ and $\mathcal{T}_n^+$.**
Let I be any ordered set, and let $G := \bigoplus_{i \in I} \mathbb{Z}$ be the group with
lexicographical order "with significance according to I". More precisely,
for $a = (a_i)_{i \in I} \in G$, set $v(a) := \max\{i \in I \mid a_i \neq 0\}$ if $a \neq 0$ and
$v(0) := -\infty$. (This is well-defined, since only finitely many a_i are
non-zero.) Now define the order on G by $a > 0$ iff $a \neq 0$ and $a_{v(a)} > 0$.

For $j \in I$, let us write g_j for the map $\mathbb{Z} \to G$ sending $\mathbb{Z}$ to the j-th
summand of G. Now let us determine $\mathcal{S}_n$ (for $n \geq 2$). For $j \in I$,
the largest convex subgroup not intersecting $g_j(1) + nG$ is $H_{<j} :=$
$\{g \in G \mid v(g) < j\}$, thus we get an injection $I \hookrightarrow \mathcal{S}_n$. For arbitrary
$a = (a_i)_{i \in I} \in G \setminus nG$, we do not get more groups: $G_{\mathfrak{s}_n(a)}$ is equal to H_j,
where $j \in I$ is the largest index such that $a_j \notin n\mathbb{Z}$. Thus $\mathcal{S}_n$ is equal
to I, possibly enlarged by one element corresponding to the group $\{0\}$
(since $G_{\mathfrak{s}_n(a)} = \{0\}$ for $a \in nG$). In particular, I can be interpreted in
G.

Now consider the sorts $\mathcal{T}_n$ and $\mathcal{T}_n^+$. The group $G_{\mathfrak{t}_n(a)}$ is the union
of all $H_{<j}$ not containing a, so it is equal to $H_{<v(a)}$; still nothing
new. However, $G_{\mathfrak{t}_n(a)+}$ is the intersection of all $H_{<j}$ containing a, i.e.
$G_{\mathfrak{t}_n(a)+} = H_{\leq v(a)} := \{g \in G \mid v(g) \leq v(a)\}$ which might be a group
which we did not have before.

Now modify our example by choosing a subset $I' \subseteq I$ and by replac-
ing, for each $j \in I'$, the factor $\mathbb{Z}$ of G by $\mathbb{Q}$. Then, $\mathcal{S}_n$ parametrizes only
those groups $H_{<j}$ for which $j \in I \setminus I'$. However, elements $a \in G$ with
$v(a) \in I'$ can still be used to obtain elements of $\mathcal{T}_n$ and $\mathcal{T}_n^+$; thus now
all three sorts can be really different. To give an extreme example, take
$I = \mathbb{R}$ and $I' = \mathbb{R} \setminus \mathbb{Q}$; then, as ordered sets, we have $\mathcal{S}_n \cong \{-\infty\} \,\dot{\cup}\, \mathbb{Q}$,
whereas $\mathcal{T}_n \cong \mathcal{T}_n^+ \cong \{-\infty\} \,\dot{\cup}\, \mathbb{R}$.

4.3. **An example for $G_\alpha^{[n]}$.** In general, the group

(*) $$H_1 := G_\alpha^{[n]} = \bigcap_{H \supsetneq G_\alpha} (H + nG)$$

(where H runs over convex subgroups of G) is not of the form $H_0 + nG$
for any convex subgroup H_0 of G. Here is an example. We use the
notation from Section 4.2. Let $I = \mathbb{N}$, but with reversed order; set
$G' := \bigoplus_{i \in I} \mathbb{Z}$ (ordered as in Section 4.2), and let G be the subgroup
of G' consisting of those $(a_i)_{i \in I} \in G'$ with $\sum_i a_i \in n\mathbb{Z}$ (for any fixed
$n \geq 2$).

Choose $\alpha := \mathfrak{s}_n(0)$ and define H_1 by (*). Then $G_\alpha = \{0\}$, and
the largest convex subgroup of G contained in H_1 is $\{0\}$, so the only
candidate of the form $H_0 + nG$ which could be equal to H_1 is nG itself.

Any element $(a_i)_{i \in I} \in nG$ satisfies $\sum_i a_i \in n^2\mathbb{Z}$. On the other hand, for any non-trivial convex subgroup $H \subseteq G$, we have $H + nG = H + nG'$, since the condition $\sum_i a_i \in n\mathbb{Z}$ can always be satisfied by adding an element of H. Thus $H_1 = nG'$, which is strictly bigger than nG.

References

1. Oleg Belegradek, Viktor Verbovskiy, and Frank O. Wagner, *Coset-minimal groups*, Ann. Pure Appl. Logic **121** (2003), no. 2-3, 113–143. MR MR1982944 (2004i:03062)
2. Raf Cluckers and Immanuel Halupczok, *Approximations and lipschitz continuity in p-adic semi-algebraic and subanalytic geometry*, 2010, Preprint.
3. Y. Gurevic, *The decision problem for some algebraic theories*, 1968, Doctor of Mathematics dissertation.
4. Y. Gurevich and Peter H. Schmitt, *The theory of ordered abelian groups does not have the independence property*, Trans. Amer. Math. Soc. **284** (1984), no. 1, 171–182. MR MR742419 (85g:03053)
5. Yu.Sh. Gurevich, *Elementary properties of ordered Abelian groups.*, Am. Math. Soc., Transl., II. Ser. **46** (1964), 165–192 (English. Russian original).
6. Immanuel Halupczok, *A language for quantifier elimination in ordered abelian groups*, Séminaire de Structures Algébriques Ordonnées 2009–2010 (Françoise Delon, Max A. Dickmann, and D. Gondard, eds.), vol. 85, Équipe de Logique Mathématique, Paris, 2011.
7. Bruno Poizat, *Cours de théorie des modèles*, Bruno Poizat, Lyon, 1985, Une introduction à la logique mathématique contemporaine. [An introduction to contemporary mathematical logic]. MR 817208 (87f:03084)
8. Matatyahu Rubin, *Theories of linear order*, Israel J. Math. **17** (1974), 392–443. MR 0349377 (50 #1871)
9. P. H. Schmitt, *Model theory of ordered abelian groups*, 1982, Habilitationsschrift.
10. Peter H. Schmitt, *Model- and substructure-complete theories of ordered abelian groups*, Models and sets (Aachen, 1983), Lecture Notes in Math., vol. 1103, Springer, Berlin, 1984, pp. 389–418. MR MR775703 (86h:03063)
11. Volker Weispfenning, *Elimination of quantifiers for certain ordered and lattice-ordered abelian groups*, Proceedings of the Model Theory Meeting (Univ. Brussels, Brussels/Univ. Mons, Mons, 1980), vol. 33, 1981, pp. 131–155. MR MR620968 (82h:03022)

Université Lille 1, Laboratoire Painlevé, CNRS - UMR 8524, Cité Scientifique, 59655 Villeneuve d'Ascq Cedex, France, and, Katholieke Universiteit Leuven, Department of Mathematics, Celestijnenlaan 200B, B-3001 Leuven, Belgium
E-mail address: `Raf.Cluckers@math.univ-lille1.fr`
URL: `http://math.univ-lille1.fr/∼cluckers`

Institut für Mathematische Logik und Grundlagenforschung, Universität Münster, Einsteinstrasse 62, 48149 Münster, Germany
E-mail address: `math@karimmi.de`
URL: `http://www.immi.karimmi.de/en.math.html`